

A technosocial epistemology of Wikipedia

Elena Falco

University College London (UCL)

Department of Science and Technology Studies

Submitted for the degree of Doctor of Philosophy

March 2024

Declaration

I, Elena Falco, confirm that the work presented in my thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

The project is composed of an empirical and a theoretical part, which inform each other iteratively. I used a mixed-methods approach for the empirical research, drawing tools from science and technology studies and platform studies. Specifically, I have conducted interviews with members of the Wikipedia community – both volunteers, and employees of Wikipedia’s publisher, the Wikimedia Foundation; analysed community pages and technical documentation published by Wikipedia or the Foundation; explored the historical and ideological connections between Wikipedia and its major intellectual predecessors: traditional encyclopaedias, Ayn Rand and the Free and Open Source Software movement. From a theoretical point of view, I have drawn concepts from empirically informed, contemporary philosophy of technology, such as postphenomenology, value sensitive design, and adjacent work.

The upshot of my thesis is twofold, and consists of some empirical findings, and theoretical contributions. The empirical findings consist in showing how values are operationalised through the design of Wikipedia’s coded infrastructure, and describing the ideological roots of Wikipedia’s design. I also conceptualise en.wikipedia.org as a site of epistemic debate, which is carried out by those who contribute code to the website, through the code they write.

From a theoretical standpoint, I propose to complicate the current story told by approaches such as values in design and value sensitive design, by highlighting the role of the artefact itself in the process of embedding. I propose to conceptualise the values embedded in design as co-constructed by the artefact itself and its designers, and to keep into account how the materiality of complex artefacts can influence trade-offs and relationships between values.

Impact statement

The work presented in this dissertation concerns one of the most accessed and influential websites in the world, as well as a widely used source of research data. As such, its impact is potentially far-reaching, crossing geographical and disciplinary boundaries. The insight obtained from studying Wikipedia is potentially generalisable to other coded artefacts, providing a framework for further study, and potentially, inspiration for developers. The academic impact of my work spans both scholarship and practice. In terms of scholarship, my study addresses a gap in the analysis of Wikipedia itself, and furthers our understanding of the ethics of technology. The process of designing and maintaining Wikipedia's platform hasn't been given extensive attention so far; here, I provide a novel account that will expand our knowledge of how Wikipedia works. The theoretical insight I obtain from my study can inform further empirical investigations of coded artefacts, especially platforms developed for the handling of knowledge, such as wikis in general, peer-review systems, coordination tools, project management systems, productivity tools, referencing systems. Some of the principles described below can be applied to coded objects in general, providing tools to understand other programmed artefacts that define our era: machine learning-based programs (AI), ICT, videogames, interfaces (software, Web and app), social media, robots, and so forth. When it comes to academic practice, it is important to note that Wikipedia is widely used as a data source, as well as an educational tool. Highlighting the structural constraints of Wikipedia might inform research methodologies, by situating data, contextualising how content is created and how interaction between members of the community is mediated. In educational settings, Wikipedia is typically either framed as a source (and as such, its use is either encouraged or discouraged) or as a tool: editing Wikipedia is a popular student activity. My work can be used to both teach how to use Wikipedia, and make decisions on its role as an educational tool. Wider social benefits of my research are two-fold. First, dissemination of my work – through the means of popular press, events presentations, participation to podcasts, and so forth – can help those who use Wikipedia understand it better, and calibrate their relationship to Wikipedia's content accordingly. Secondly, the generalisable insight developed by using Wikipedia as a case study can be used by industry professionals – designers, developers, decision-makers – who develop technology within an ethical framework to improve their methodologies. Should circumstances allow, I would be keen to work with industry professionals, on co-development projects, in which case the impact would be tangible: one, or more, material artefacts built with the aid of my research.

Acknowledgments

During my PhD, I discovered a book called *Wayward Lives, Beautiful Experiments*, by the dazzling Saidiya Hartman. It's a lovely book, but that's not the point here. The point is: its title perfectly describes the last five years. Pursuing a doctorate has been a profoundly transformative experience. A beautiful experiment – personally, professionally, and in the muddy middle where those categories mingle (also known as: me). Above anything else, the last five years have taught me the value, the bravery, the joy, of waywardness.

In that spirit, I would like to thank, first, the Alice Farrands room dwellers. I can say with confidence that, if I hadn't met you all, my life, my thesis, my career, would be extremely, and regrettably, different. Some special mentions are in order. Kylo Thomas, for seeing me; for provoking me; for causing a nuisance with me. Jas Crowson, for being my sister in weirdness. Honorary member Stephen Hughes, for always being game. Ben Weil, for welcoming me on my first day, prompting me to do the same with others.

I would also like to thank my supervisors, and those who have mentored me, in big and small ways. Phyllis Illari, for modelling softness and might. Tiago Mata, for being there across borders, two children, and a pandemic. Federica Russo, for making me feel like I belong. Rune Nyrup, for encouraging me to stand up and speak up. Emily Dawson, for her candour. Emma Tobin, for taking me seriously. Chiara Ambrosio, for a staggering number of Moretti. Cian Donovan, for picking my brain.

I would like to thank the Wikipedians who shared their experience and views with me – and those who refused to do so. Buildings are made out of solids and voids.

Nothing happens in a vacuum, though. Especially not doctorates. I am where I am because of a rich and generous community. My deepest, most loving gratitude goes to Ian Bridgeman, for possibly the best off-the-cuff pep talk ever delivered on Venice Beach; for his valuable insight as a nerd, Web developer, and writer; and for always bringing me coffee. My mum, Nunzia Curci, for her intelligence, drive and grit. My dad, Beppe Falco, for always telling me I had taken after my mum (and, secondarily, letting me stick around while he disassembled computers). All my friends, for pretending they have finally understood what STS is. Davide Pizzeghello and Beppe Iannello, for bringing laptops to pubs long before it was acceptable, and worse yet: before they had Wi-Fi.

I wish to dedicate this thesis to my grandfather, Aniello – Nonno Nello for me, Mast'Aniello for his peers. Maker, migrant, polyglot, chequers player, opera nerd.

Quiet teacher.

Loud anti-fascist.

Table of contents

Declaration 2

Abstract 3

Impact statement 4

Acknowledgments 5

Table of contents 6

Style note 10

Introduction and structure of the thesis 11

Research questions 12

Roadmap 12

1. An introduction to Wikipedia 14

1.1. *Introduction* 14

1.2. *A note on sources* 15

1.3. *How Wikipedia works, in practice* 16

1.3.1. The editing process 16

1.3.2. Wikipedia's policies 20

1.3.3. Coded objects 25

1.4. *How Wikipedia works, in theory* 27

1.4.1. Wikipedian epistemologies of Wikipedia 27

1.4.2. Independent, scholarly epistemologies of Wikipedia 28

1.5. *Wikipedia's roots* 30

1.5.1. Historical overview 30

1.5.2. Hacker culture and the free software movement 32

1.5.3. Ayn Rand 36

1.5.4. Encyclopaedias 38

1.5.5. Design choices 42

Conclusion 43

2. Methodology and theoretical framework 44

2.1. *Introduction* 44

2.2. *Research questions* 44

2.3. *Theoretical framework* 47

2.3.1. Postphenomenology 47

2.3.2. Value-led design theories 52

2.3.3.	Platform studies	55
2.4.	<i>Data collection</i>	57
2.4.1.	Approaching Wikipedia	57
2.4.2.	Conversations with Wikipedians	59
2.4.3.	Documentary sources	60
2.5.	<i>Data analysis</i>	62
2.5.1.	Early work on talk pages	62
2.5.2.	Thematic analysis	63
2.5.3.	Platform analysis	64
2.5.4.	Limitations	65
2.6.	<i>Access</i>	66
2.6.1.	Background	66
2.6.2.	Recruitment of participants	68
2.6.3.	The politics of access	69
2.6.4.	Access to material sources	70
2.6.5.	Ethics	71
	<i>Conclusion</i>	73
3.	Wikipedia's aesthetic	74
3.1.	<i>Introduction</i>	74
3.2.	<i>Interface design and Wikipedia's sustainability</i>	74
3.2.1.	Accessibility	75
3.2.2.	Engagement	77
3.3.	<i>Abstraction, situatedness and objectivity</i>	82
3.3.1.	Abstraction	82
3.3.2.	Situatedness	93
3.3.3.	Abstraction, situatedness & justice	95
3.4.	<i>Belonging</i>	99
3.5.	<i>Modernism and anxieties around progress</i>	105
3.6.	<i>Geek culture, antiaesthetic and antisensuality</i>	111
	<i>Conclusion</i>	115
4.	Programming an encyclopaedia	117
4.1.	<i>Introduction</i>	117
4.2.	<i>Coding as means of knowledge-production</i>	117
4.2.1.	Writing as, and, through coding	118
4.2.2.	Validation through code	119
4.2.3.	Curation with and through coding	124
4.2.4.	With or through?	125
4.3.	<i>Sorting technologies</i>	128
4.3.1.	Explicit sorting	128
4.3.2.	Sorting through affordance: litmus technologies	130
4.3.3.	Material sorting: wikitext and angle brackets	137
4.3.4.	Coded custodians	138

4.4.	<i>The construction of the knowing subject on Wikipedia</i>	140
4.4.1.	Bracketing the body: the clean channel ideal	140
4.4.2.	Celebrating bodies: userboxes and the Teahouse	142
4.4.3.	Sorting through bodies: operationalised trust and positionality	143
4.5.	<i>Negotiating epistemology through coding</i>	145
4.6.	<i>Technoepistemic practices</i>	147
4.7.	<i>Poietic objectivity</i>	150
	<i>Conclusion</i>	152
5.	Towards a material, relational account of the value-ladenness of technology	153
5.1.	<i>Introduction</i>	153
5.2.	<i>Wikipedia's value system</i>	154
5.2.1.	Neutrality	154
5.2.2.	Equity	154
5.2.3.	Inclusivity	154
5.2.4.	Structural openness	155
5.2.5.	Radical collaboration	155
5.3.	<i>Values in technology and hylomorphism</i>	156
5.4.	<i>Challenges to the hylomorphic model, from my work on Wikipedia</i>	158
5.4.1.	Materiality and the subjectivity of designers	158
5.4.2.	Material affordances and value embeddings	159
5.4.3.	Material trade-offs	162
5.4.4.	Value change	163
5.4.5.	Material legacy	163
5.5.	<i>Co-constitution of subject and object in the design process</i>	165
5.6.	<i>Co-construction of values</i>	167
5.6.1.	Selection	168
5.6.2.	Construction	169
5.7.	<i>Material dynamics</i>	170
5.7.1.	Internal dynamics	171
5.7.2.	Lineage	172
	<i>Conclusion</i>	173
	Conclusion	175
	Bibliography	177
	List of cited pages from Wikipedia and related wikis	197
	Appendix A: Epistemic culture themes	203
	Appendix B: Infrastructure and design themes	206
	Appendix C: Codebook	208

Style note

A note on citing Wikipedia: when citing wikis, the APA style, which I use throughout this project, suggests using the title of the article, surrounded by quotes, and the year of retrieval. This method is unsuitable for me, for two reasons: first, it doesn't allow for pages that are not encyclopaedic articles, such as information pages. Secondly, Wikipedia, in the context of this dissertation, is to be treated mostly as a primary source (more in section 1.2). Hence, I devised my own convention: I will cite Wikipedia pages by including, in the in-text citation, the title as it appears on Wikipedia, or in absence of a title, the last section of the url, surrounded by quotes, or, when in parenthesis, without quotes.

I have decided to adopt a similar notation for wikis related to Wikipedia, which concern themselves with operational matters. These are MediaWiki, the wiki containing technical documentation about, and coordinating work around, Wikipedia's software (<https://www.mediawiki.org/wiki/MediaWiki>); Meta-Wiki, used to coordinate, and publish information about, Wikimedia Foundation projects (https://meta.wikimedia.org/wiki/Main_Page); and the Wikimedia Foundation Governance Wiki (<https://foundation.wikimedia.org/wiki/Home>). Pages from the wikis above will be cited with their titles as they appear on the Web page, preceded by, respectively, "MediaWiki:", "MetaWiki:" and "Foundation:".

Wikipedia-related references are listed separately from the main Bibliography, in the section "List of cited pages from Wikipedia and related wikis".

Introduction and structure of the thesis

I became interested in the technical infrastructure underlying knowledge-production while working in the backstage of academic publishing, as one of those obscure professionals whose job is answering emails of frustrated researchers trying to change the order of authors on a submission, sending reminders to overworked reviewers, and patient updates to appease corresponding authors of papers under review (often, and with surprisingly little effect on reflexivity, to individuals fulfilling both roles). In those days, my main work companion, gripe, and support was a piece of software called Editorial Manager. Editorial Manager is a layered mountain – a byzantine composite of pieces of software created and deployed at different times, by different people, with the technologies available at the time of programming. Navigating Editorial Manager was like walking around a fairground castle – a nervous trod through a series of jarringly linked, heterogeneous rooms populated by a cast of grotesque characters.

And that, I used to muse, was the backbone of truth. Our society rests on a deep, desperate trust in science. We look at peer-review as a beacon of hope – the process of peer-review being an essential component of our faith in the existence of an independent, knowable reality. Yet it rests on fiddly interfaces, overworked academics, and unacknowledged administrative staff.

Another topic I used to think about back then was the ethics of peer-review. One example above all: not many know that some submissions are desk rejected before they even reach an academic editor, for issues that have to do with their formatting, or, under the euphemistic guise of “beyond scope of the journal”, common sense. I have personally desk rejected countless submissions connecting vague claims about the energy of the universe, quantum physics, and the power of the mind. I have also rejected articles because the email addresses of their corresponding authors were not institutional email addresses. Now, there are parts of the world where institutional email addresses don’t exist; and what of independent research? The practice of filtering email addresses is heavily loaded, and at the time, no one, despite my attempts to raise the issue, seemed to be interested in thinking about its ethical fallout.

The question of how, and in what measure, processes of validation of knowledge influence what we think of as the truth, is the pulsating core of my doctoral project. Wikipedia is an interesting focus because, as well as being possibly the most-read encyclopaedia in the world, it is also a gigantic machine for the vetting – or, as most would agree on, creation – of knowledge. Wikipedia is both an editorial and a technological enterprise. As well as asking how the process of content-production influences content itself, we need to ask ourselves how the converse: does the editorial line of the encyclopaedia influence its infrastructure? What is the shape of the relationship between what is written on Wikipedia and how it’s produced?

What’s more, Wikipedia emerged from a broth of cultural influences that happen to be paramount to understanding today’s technological landscape. The project was born from the free and open source software (FOSS) movement, steeped in hacker and nerd culture; Wikipedia’s

founders were sympathizers of Ayn Rand’s philosophy. Both FOSS culture and Ayn Rand are major influences on Silicon Valley, the cradle of 21st century technological thinking. Is there a connection between technological culture and Wikipedia? And if there is, does it trickle down from infrastructure to content? How? Who is most powerful – code or content?

Simmering on the backburner are, as they were during my time in publishing, troubling ethical worries. Wikipedia is a very Western product. It is a proud child of the enlightenment; it categorises knowledge in ways that will be familiar to anyone with a formal education, but which bulldoze over non-Western ontologies, much like the straight borders dividing former colonies from one another, oblivious of the terrain they are fictionalised to be traced upon. Gender discrimination on Wikipedia is also well documented (Ford & Wajcman, 2017), and reflected in its content. What is, one might ask, the link between the process – people, machinery and ideology – and the way the world is represented on Wikipedia?

As I analysed my data, I noticed that insight from Wikipedia can be fruitfully applied to other artefacts as well, in the form of a theory of technological value embeddings. Some initial thoughts on the matter are presented here as well.

The focus of my research project shifted over time, within the landscape traced above – I will provide details in 2.2.2. Here, I will simply list the research questions addressed by this dissertation, and provide a roadmap for the thesis itself.

Research questions

1. How is Wikipedia’s epistemology embedded in the design of Wikipedia’s website?
2. How have the major ideological influences on Wikipedia’s culture impacted the design of the knowledge-production process on Wikipedia?
3. What theoretical contribution to the field of philosophy of technology can be gained from a close reading of en.wikipedia.org?

Roadmap

Due to the interdisciplinary nature of my project, and to its subject, the structure of this thesis is unusual. While this is a dissertation drawing on my own empirical findings, there won’t be a chapter presenting results. Rather, I will use insight from data analysis to build an argument about the values embedded in Wikipedia’s infrastructure, presenting empirical evidence as and when needed. Technical details about data analysis are presented in appendices. Further, Wikipedia presents idiosyncratic local norms, which, to some extent, also define the notion of norm in a locally specific way, and a social structure that doesn’t mirror any arrangement already familiar to social scientists. Wikipedia’s uniqueness has a bearing on how Wikipedia can and should be investigated, and, consequently, heavily shaped my methodology. Hence, the first chapter won’t be about methodology, but about Wikipedia itself.

In Chapter 1, *An introduction to Wikipedia*, I provide an overview of Wikipedia’s community regulation and structure, an introduction to its coded infrastructure, and summarise some

influential writing published about Wikipedia's epistemology so far, by both Wikipedian and independent scholars.

In Chapter 2, *Methodology and theoretical framework*, I outline the methodology of the study, and the theoretical framework I move within. Methodology and theory are connected: I use methods frequently used in value sensitive design and leverage theory from various approaches to values in design, combined with postphenomenology.

In Chapter 3, *Wikipedia's aesthetic*, I analyse Wikipedia's aesthetic qualities, and how they relate to the production of knowledge. I analyse Wikipedia's visual appearance, writing style, formats, and show their connection to Wikipedia's internal culture, as well as external influences (chiefly hacker culture and modernist design).

In Chapter 4, *Programming an encyclopaedia*, I examine some of the technologies deployed for the production of content, their relation with content itself, and what values they embed and uphold. I then connect the making of coded infrastructure with knowledge production through the figure of the nerd, and sketch a form of objectivity that I call "poietic objectivity".

In Chapter 5, *Towards a material, relational account of the value-ladenness of technology*, I bring to bear conclusions I reached about Wikipedia, in order to critique current understandings of the value-ladenness of technology. With the support of the work of Gilbert Simondon, I argue that an account of value-ladenness in digital objects – and perhaps technology more broadly – would benefit from a closer examination of materiality. I argue that values are not embedded in technology – rather, they are co-constructed during the design process; similarly, the subjectivities of designers, often framed as the main impetus for value-embeddings, are also shaped by the practice of design.

1. An introduction to Wikipedia

There was but one course left, therefore,—to try the whole thing anew upon a better plan, and to commence a total reconstruction of sciences, arts, and all human knowledge, raised upon the proper foundations. And this, though in the project and undertaking it may seem a thing infinite and beyond the powers of man, yet when it comes to be dealt with it will be found sound and sober.

Francis Bacon, Proœmium of the Instauratio Magna

1.1. Introduction

Wikipedia describes itself as a “multilingual, web-based, free-content encyclopedia project supported by the Wikimedia Foundation and based on a model of openly editable content” (Wikipedia: About). Wikipedia’s tagline is “the free encyclopedia that anyone can edit” (Main Page). It expresses the project’s guiding principles of exhaustiveness, open access, and wide participation: the encyclopaedic character of the project implies an ambition to cover all available knowledge, in writing; anyone with Internet access can, in principle, benefit from, and contribute to produce, its content.

In this project, I will focus on the technologies that sustain Wikipedia, focusing on how they are designed, and how they intersect with Wikipedia’s epistemology. This project’s research questions and methodology have been, in part, shaped by my understanding of Wikipedia’s culture, infrastructure and power relations. The story I tell in this dissertation relies heavily on small details: how Wikipedia is edited, by whom, and in what ways; the historical development of certain specific features, and what kind of perceived need they were designed to fulfil; the subtle lattice of power relations embedded in coded infrastructure. This chapter lays out the groundwork, by leveraging existing secondary literature, primary sources, and my own data analysis¹.

The necessity of this account stems from a gap in the existing analyses of Wikipedia, focused mostly on its culture, as separate from the design of the website, or on the interactions between Wikipedia’s editors and the coded infrastructure they use to edit the encyclopaedia. As I will show in this chapter, such accounts, as well as stopping short of engaging with how the website is designed, ignore the fundamental fact that Wikipedia was an offshoot of coding culture, and absorbed, in ways that I will document over the course of this thesis, some of its culture.

Programming practice, I will argue in this dissertation, is not incidental to Wikipedia: it is one of its key components, reverberating beyond creating the necessary tools to update the website. Programming practice contributes to define the local understanding of truth, knowledge and knowers, thereby constituting an integral part of Wikipedia’s epistemology. Wikipedia’s

¹ I have conducted two thematic analyses over the course of this project, for reasons that will be detailed in chapter 2. Details of the two thematic analyses can be found in the appendices: *Epistemic culture* (Appendix A) and *Infrastructure and design* (Appendix B).

organisational infrastructure, I will argue, is entangled with, and contributes to constituting, Wikipedia's ideology, alongside external influences such as Ayn Rand's thought, modernist ideals of progress, and hacker culture.

1.2. A note on sources

As well as drawing on my own data analysis, this chapter uses a mix of primary and secondary sources, in addition to a third group of sources that sits between these two categories. I will briefly explain what this means and why it matters.

Primary sources used for this chapter include interviews, Wikipedia policies, community pages, first-hand accounts of events, and a manual for the use of Wikipedia written by Wikipedians themselves. Secondary sources include ethnographies of Wikipedia, and scholarly analyses of Wikipedia in terms of its politics, epistemology, and so forth. The third group evades categorisation, because it is populated by works that either don't technically qualify as secondary sources, but I will use them as such in some cases, or do technically qualify as secondary sources, but I will treat them, occasionally, as primary.

The former – primary sources that I will use, occasionally, simply as sources of information, comprises factual Wikipedia pages about Wikipedia, and *Wikipedia: The Missing Manual* (Broughton, 2008). Wikipedia's history is not widely documented, especially not outside of the community (more on this below). Hence, when it comes to relatively inconsequential facts with no apparent ideological bearing, I will use Wikipedia's own account of its own history. There are exceptions to this: for instance, I found out during data collection that the history of Wikipedia's logo is misrepresented on Wikipedia for reasons that will become clear in chapter 3. Equally, I will take basic procedural documents as factual. When presenting policies, however, I will not take them at face value: rather, I will provide context to facilitate their interpretation.

The latter – secondary sources that I will handle carefully – comprise writing about Wikipedia written by authors who are Wikipedians themselves. For instance, *The Wikipedia Revolution: How a Bunch of Nobodies Created the World's Greatest Encyclopedia*, by Andrew Lih (Lih, 2009), is a book about early Wikipedia history, and contains precious accounts of early events that shaped the project. However, as may be gleaned from the title, Lih is an enthusiastic Wikipedian: based on his LinkedIn profile, Lih is currently "Wikimedian at Large" at the Smithsonian Institute, and "Wikimedia Strategist" at the Met Museum; Jimmy Wales, founder of Wikipedia, wrote the foreword to the book. Lih's book is not an isolated case: the bulk of the factual literature about Wikipedia is written by people who have ties to Wikipedia, such as Dariusz Jemielniak, who is a Wikipedia editor and a member of the Wikimedia Foundation Board of Trustees, and Joseph M. Reagle, who based on his CV, participated to Wikimania, Wikipedia's annual meetup, in 2006, before he started writing about Wikipedia. Jemielniak and Reagle are academics: their work is, technically, scholarly; at the same time, they don't hide their partisanship: while acknowledging Wikipedia's shortcomings, they write about it, mostly, to celebrate it. Hence, while these are precious sources for historical, ethnographic and procedural matters, I use them critically, much like I would primary sources.

1.3. How Wikipedia works, in practice

It is commonplace, within the Wikipedia community, to joke that Wikipedia works in practice, but not in theory. Wikipedia is clearly a successful project – a large, and broadly speaking reliable online encyclopaedia, available for free and in several languages – but, the narrative around the joke goes, it’s unclear how that is possible, given that it is made by volunteers, who are also non-experts in the fields they cover and with a loose (some would say: non-existent) organisational structure. It has been indeed a challenge for social epistemologists and philosophers to understand how, as the subtitle of one Wikipedia’s histories goes, “a bunch of misfits built the world’s greatest encyclopaedia” (Lih, 2009).

Without necessarily buying into hyperbole, the aim of the present work is, indeed, to contribute to a theoretical understanding of how Wikipedia works, in theory. First, however, I will lay out my understanding of how Wikipedia works in practice: that is, the procedures, rules, social arrangements, and tools involved in creating and editing Wikipedia’s content. The following account is built on direct review of a range of primary sources, such as policies, guidelines, and community pages, as well as secondary and pseudo secondary sources (see section 1.2 for details). While aiming to as factual an account as possible, a degree of analysis will be necessary, as certain parts of Wikipedia’s culture, such as, for instance, policies, are difficult to understand when taken at face value. The very nature of policies, as understood within the Wikipedia community, is dissimilar to how policies are understood in most organisations, calling for some contextualisation. I will then, in the following section, sketch the landscape of theoretical accounts of Wikipedia, including reflexive work by Wikipedians themselves.

1.3.1. The editing process

At the time of writing, English-language Wikipedia contains 6,753,769 content articles (Wikipedia: About). Each article consists of an article page, and a talk page. The article page contains the content of the article, while the talk page is dedicated to discussing changes to the article. Each – Article and talk pages – feature three tabs: Read, Edit, and View history. The Read tab is fairly straightforward, and the default view. The Edit allows allows to edit the page: in the case of the Article, editing means editing the content of the article; in the case of the talk page, editing allows to participate to the discussion, by inserting one’s post in the feed. The View history tab allows to see the list of edits to the page from its inception, including data such as timestamp and author, and compare versions of the page.

Wikipedia is edited by volunteers, who decide autonomously what to edit and how, in accordance with policies (described below). Editing a Wikipedia article consists, in practice, in clicking on the Edit tab of the article page, and making changes to the editing interface. Once changes are made, they can be committed to the article page. Changes can be reverted: editors can click on the “undo” button near any edit, thereby restoring the previous version of the article. A principle known as “radical collaboration” underlies Wikipedians’ ability to undo each other’s work: no one’s work is, in principle, untouchable. It is considered good practice to append an edit summary to each edit, i.e. an explanation of why the edit was made in the first place, or what kind of contribution it makes to the page as a whole. In essence, editing Wikipedia is a process of constant addition and

subtraction of small portions of content (mostly text, as well as images, data and so forth).

Reverting someone's edit signifies, in practice, disagreement: it means the editor who reverts thinks the edit was unwarranted in the first place. It used to be possible to revert other people's edits indefinitely, leading to a phenomenon known as "edit warring" (Wikipedia: Edit warring). Eventually a cap on edits was introduced: after three reverts on the same edit, the disagreement should be settled by discussing on the talk page. Once a disagreement reaches the talk page, anyone can weigh in on either side. It is not unusual for conversations on talk pages to sprawl far beyond the initial issue to be settled. So much so that each talk page starts with a fixed yellow box carrying the disclaimer "This is for general discussion of the article's subject".

According to accounts written by Wikipedians, the Wikipedia community is run in an egalitarian spirit (Ayers *et al.*, 2008; Jemielniak, 2014). However, "parahierarchies" (Jemielniak, 2014) are present: in practice, leaders exist but their power is only partially, and informally, acknowledged. Users can be ranked by levels of access, that is by the power to make changes to infrastructure, granted to them in function of their acquired status within Wikipedia. Status within Wikipedia is gained by accruing large amounts of edits, and editing in a way that the community finds helpful. External status is not explicitly acknowledged, and editors are assumed to start on equal footing when they join the community. The groups are: users, unregistered users, registered users, autoconfirmed and confirmed accounts, administrators (sysops), bureaucrats, stewards, founder (singular: this group includes exclusively Jimmy Wales).

Unregistered users don't have a username, and can be recognised through their IP address. Their powers are limited: they can only edit pages which are not protected or semi-protected (Meta: Unregistered user). Discrimination is practiced against unregistered users, with recurring calls to deprive them of editing rights, raised since at least 2004 (Meta: Anonymous users should not be allowed to edit articles; Wikipedia: Editors should be logged-in users (failed proposal)). An argument often used against unregistered users is that of vandalism: a figure quoted often is that 80% of vandalism comes from unregistered users (Wikipedia: IP addresses are not people; Meta: Musings about unregistered contributors). Pushback against anti-anon sentiment has taken the form of direct defences, policies, and defence of unregistered as a legitimate choice, some kind of identity (Wikipedia: IPs are human too; Meta-Wiki (Musings about unregistered contributors; Wikipedia: Welcome unregistered editing; Meta: Association of Good Faith Wikipedians Who Remain Unregistered on Principle).

Registered users are those users who have signed up with a username, have a few more rights than unregistered users, including, after "a few days" of editing articles, creating a new page (Wikipedia: Your first article). The exact conditions by which a registered user becomes autoconfirmed (hence able to start a new page) are somewhat vague (Wikipedia: User access levels). It is a factor of time since registration and number of edits made (*ibid.*). Confirmed users are users who are granted the same privileges as autoconfirmed, before they are automatically recognised as sufficiently experienced.

Administrators, as well as enjoying the rights of registered users, can block users and IPs, edit protected pages, override blacklists, and give registered users ad hoc privileges, which they can also assign to themselves, should they see the need (Wikipedia: User access levels). Administrators' power is constantly downplayed. Becoming an admin is known as "taking the mop" (Jemielniak, 2014), as

they are supposedly carrying out janitorial duties, keeping the community in good shape. In fact, the symbol of administrators is the Wikipedia logo with a mop in front of it. Jimmy Wales himself has weighed on the issue:

I just wanted to say that becoming a sysop is *not a big deal*.

I think perhaps I'll go through semi-willy-nilly and make a bunch of people who have been around for awhile sysops. I want to dispel the aura of "authority" around the position. It's merely a technical matter that the powers given to sysops are not given out to everyone.

I don't like that there's the apparent feeling here that being granted sysop status is a really special thing.

(Jimmy Wales, cited in Jemielniak, 2014)

The first interesting thing about this quote is that “semi” before “willy-nilly” and the qualification “who have been around for awhile”: if Jimmy Wales is not prepared to make anyone administrator, it means that the exclusion of some people from being sysops is not “merely a technical matter”.

Additionally, the “aura of ‘authority’” mentioned by Wales is, analytically, a function of the perception of administrators within the community; if they are seen as powerful, they are powerful. Administrators are the lowest level of users with power over other users: it is unsurprising that it would be culturally associated with status. Not just that: in order to become an administrator, one needs to be elected by public vote. Even unregistered users can vote for this role (while of course they can't be elected). While any registered user can submit a Request for Adminship (RfA), only users who are well-known to the community are likely to receive a number of endorsements sufficient for election (Jemielniak, 2014; Wikipedia: Requests for adminship). The process of election includes also a harsh vetting process which includes argumentation with other users (Jemielniak, 2014). Questions are asked and the potential administrator has to defend themselves. The harshness of this process is defended as preparation for future work, when the administrator will have to engage with argumentative users (*ibid.*). With 36,124,752 registered users (14 April 2019), of which 139,560 active (14 April 2019), only 1,179 are administrators (Meta: Special: Statistics).

Bureaucrats are administrators who can remove user rights, including those of admins and bureaucrats themselves. They are also expected to resolve disputes “in difficult cases” and explain their decision (Wikipedia: Requests for adminship). There are currently 21 bureaucrats in the English Wikipedia. The election process is similar to that of administrators, but with stricter thresholds: consensus from the community needs to reach around 85% (Wikipedia: Requests for adminship). Stewards are active on a global scale (while administrators and bureaucrats have only local jurisdiction). They have the same user rights as administrators and bureaucrats, but on a global scale. (Meta: Stewards). They are supposed to use their powers only in case of emergency (*ibid.*), and are not expected to weigh in on disputes, that are supposed to be solved at the local level (*ibid.*). They are elected annually in a global election (*ibid.*). Finally, Jimmy Wales, Wikipedia's founder or co-founder, enjoys his own access level, known as “founder”. He has “full access to user rights” (Wikipedia: User access levels). His leadership evolved over time: in the early days of Wikipedia, his

authority was taken as a given (Jemielniak, 2014), and he was able to write pieces such as the “Statement of principles”, whose introduction states.

As we move forward with software and social changes, I think it is imperative that I state clearly and forcefully my views on openness and the licenses. This page, like all Wikipedia pages, is a living, dynamic document, which I will update and clarify as legitimate questions arise.

I should point out that these are my principles, such that I am the final judge of them. This does not mean that I will not listen to you, but it does mean that at some ultimate, fundamental level, this is how Wikipedia will be run.

(User: Jimbo_Wales/Statement of principles)

Over time, his leadership style went from interventionist to more relaxed (Jemielniak, 2014). In fact, he has been known to not use his founder powers, in order to avoid criticism from the rest of the community (*ibid.*). Nowadays his role is more that of spokesman and inspiration to the community (*ibid.*), interacting with the community through Wikipedia and at events, “like a prophet meeting his flock” (Ayers *et al.*, 2008, p. 8).

Hindrances to the regular functioning of Wikipedia are known as “disruptive editing” (Wikipedia: Disruptive editing). Disruptive editing is a somewhat vague concept; the examples outlined on (Wikipedia: Disruptive editing) gesture towards a general disrespect of community practice, which makes it more difficult for well-intentioned editors to do their job. Listed behaviours encompass continuous breach of policies, being uncooperative in debate, misusing the citing facilities in order to make a point (such as adding “citation needed” to all sentences in a page, even when those are present, to show that the whole page should be rewritten or deleted).

A more dangerous subset of disruptive editing is vandalism, defined on (Wikipedia: Vandalism) as “editing (or other behavior) deliberately intended to obstruct or defeat the project’s purpose”. Breaches range in severity, from removing good quality content without backing up the decision with reasons, to including blatantly irrelevant information or offensive content within an existing page. Edits that constitute clear acts of vandalism should be reverted, and vandals may be banned from the site.

A less antagonistic practice, which is nevertheless contrary to the declared aims of Wikipedia, is that of “conflict-of-interest editing” (Conflict-of-interest editing on Wikipedia), that is, editing done to advance a personal or organisational agenda. This is normally done through “sockpuppetry”, that is assuming multiple online identities, with malicious intent. Wikipedia only allows one username per person, with few exceptions. Conflict-of-interest editing has been done for political and PR reasons, to such an extent that a major sockpuppet investigation was launched in 2012 (Conflict-of-interest editing on Wikipedia). After some tension with PR firms, Wikipedia is now allowing PR edits, as long as the conflict of interest is disclosed in the talk page (Cipr, 2014; Wikipedia: Conflict of interest). Purposeful editing is leveraged by the community to support specific aims, such including more biographies of women, or expand the coverage of underrepresented areas and communities, such as Africa (Meta: Wiki Loves Africa) and women (Wikipedia: WikiProject Women in Red).

The ultimate punishment for misbehaviour is blocking, which consists in temporarily or

indefinitely suspending editing privileges, for behavioural reasons: blocked users might have acted disruptively, or in a way that threatens to damage Wikipedia or the Wikimedia Foundation. (Wikipedia: Blocking policy). A language similar to that commonly applied to criminals is often used when describing blocked users: a block is “enforced”, circumventing a block is known as “evasion”; blocks should be “preventative”, not “punitive” (*ibid.*).

Explaining the reasons for which a user may be blocked for breaching a code of conduct also forces policy to define that code of conduct. Many articles have been written about how to behave on Wikipedia, but perhaps the most concise definition of such a code of conduct is found on the page on blocking: a user is blocked when “when his or her conduct is inconsistent with a civil, collegial atmosphere and interferes with the process of editors working together harmoniously to create an encyclopedia” (*ibid.*).

As well as giving a general idea of how Wikipedia works, I have outlined access levels to show how power is distributed and obtained on Wikipedia, an issue that will become important over the course of this work. It is also important to emphasise here how the Wikipedia community deploys mechanisms that allow it to defend itself from outsiders, for instance in the form of vandalism. Wikipedia’s openness is not absolute: boundaries of belonging are sometimes demarcated clearly, as it happens, for instance, through access levels. Other times, belonging is modulated in subtler ways: in chapter 4, I will detail how technology and design function as mechanisms of boundary maintenance; in the next section, I will show how Wikipedia’s policies fulfil an analogous role, by expressing Wikipedia’s ideology, thereby attracting those who agree with it, and creating a social environment that is hospitable for some, but not others.

1.3.2. Wikipedia’s policies

As outlined in section 1.3.1, Wikipedia is edited through a collaborative, decentralised process. While, in principle, anyone can edit Wikipedia, and the editing process is, in essence, impersonal, in practice different levels of participation are afforded to different people, and editing involves interaction with a complex social world. Additionally, while editors are free to edit whatever page they have editing access to, and however they see fit – edits are committed to the page without being reviewed; reverting occurs after publication – in practice not just anything remains on a Wikipedia page. The collective process of editing Wikipedia is regulated by a complex system of policies and guidelines, which, in turn, hide an undergrowth of unspoken social rules, some of which apply to how policies and guidelines are used in the first place. In this section, I will outline Wikipedia’s main policies, and their role within Wikipedia’s culture: policies within this context are not, as one might imagine at first brush, rules that can be consistently and reliably upheld in order to sanction infractions. In other words, they are policies only in name. In practice, they function differently depending on their audience, inviting or inhibiting action; they are means by which Wikipedia’s ideology is expressed and preserved.

Wikipedia’s main policies are called the “Five Pillars” (Wikipedia: Five pillars): Neutral Point of View (NPOV), No Original Research (NOR), Verifiability (V), Ignore All Rules, and Civility. The first three – NPOV, NOR, V – concern the production of the encyclopaedia’s content; policies of this kind are known within the community as ‘content policies’ (Wikipedia: Policies and guidelines). NPOV, NOR and V are also known as ‘core policies’ (Wikipedia: Neutral point of view),

signalling their prominence within the Five Pillars.

Neutral Point of View (NPOV) established neutrality as a fundamental value. NPOV is formulated as:

All encyclopedic content on Wikipedia must be written from a neutral point of view (NPOV), which means representing fairly, proportionately, and, as far as possible, without editorial bias, all the significant views that have been published by reliable sources on a topic.

(Wikipedia: Neutral point of view)

NPOV is about the presentation of content: information needs to be laid out in a neutral manner. When a controversy is present in the sources, Wikipedians are not supposed to take sides – rather, they should represent fairly every point of view of significance. This principle applies to “both what you say and how you say it” (Wikipedia: Neutral point of view). Avoiding bias is, on one hand, a matter of not giving “undue weight” (Wikipedia: Neutral point of view) to marginal views, that are not held by a significant amount of people, and conversely downplaying mainstream views. Neutrality also applies to tone: as far as possible, Wikipedians ought to use plain language, which does not let the writers’ opinion of the topic at hand transpire.

According to Wikipedia itself, NPOV is a fundamental policy, that can’t be changed. The page that describes NPOV is protected, as shown by the lock symbol on the top-right corner of the page. Protected pages are exceptions to the “anyone can edit” policy: they can only be changed by administrators (Wikipedia: Protection policy). A further protection is advertised on a disclaimer at the top of the page: “This policy is non-negotiable, and the principles upon which it is based cannot be superseded by other policies or guidelines, nor by editor consensus” (Wikipedia: Neutral point of view). Based on NPOV’s constitutional role, it is fair to say that neutrality is a foundational value of Wikipedia.

The two other core policies, Verifiability (V) and No Original Research (NOR) are corollaries of NPOV:

[The core policies] jointly determine the type and quality of material that is acceptable in Wikipedia articles, and, because they work in harmony, they should not be interpreted in isolation from one another. Editors are strongly encouraged to familiarize themselves with all three.

(Wikipedia: Neutral point of view)

Verifiability is formulated as follows:

In the English Wikipedia, verifiability means other people using the encyclopedia can check that the information comes from a reliable source. Wikipedia does not publish original research. Its content is determined by

previously published information rather than the beliefs or experiences of editors. Even if you're sure something is true, it must be verifiable before you can add it.

If reliable sources disagree, then maintain a neutral point of view and present what the various sources say, giving each side its due weight.

All material in Wikipedia mainspace, including everything in articles, lists and captions, must be verifiable. All quotations, and any material whose verifiability has been challenged or is likely to be challenged, must include an inline citation that directly supports the material. Any material that needs a source but does not have one may be removed. Please immediately remove contentious material about living people that is unsourced or poorly sourced.

(Wikipedia: Verifiability)

In sum, Verifiability circumscribes how knowledge is justified on Wikipedia and what can be included in a Wikipedia article: only what can be traced back to reliable sources counts as knowledge, and as such, can be included on Wikipedia. Direct, personal experience doesn't count. This implication of V is more apparent in the policy's original formulation: "the threshold for inclusion is verifiability, not truth" (Wikipedia: Verifiability, original emphasis). By excluding personally held knowledge and experience, individual points of view are kept at bay, in favour of the points of view recorded in the literature, that are supposed to be, then, given the NPOV treatment – they should be represented based on relative weight in the literature.

The third core policy is No Original Research, according to which:

Wikipedia articles must not contain original research. The phrase 'original research' (OR) is used on Wikipedia to refer to material—such as facts, allegations, and ideas—for which no reliable, published sources exist.

(Wikipedia: Verifiability, original emphasis)

NOR is meant to be an additional tool in preventing bias from entering the system: on the one hand, it guarantees that information comes from a reliable source, as opposed to being the result of the primary research carried out by non-experts; on the other, it discourages experts from passing their own perspective on any given subject as uncontroversial. No Original Research is about an implied lineage, rather than the actual statement of the source: "The prohibition against OR means that all material added to articles must be attributable to a reliable, published source, even if not actually attributed" (Wikipedia: No original research).

Ignore All Rules, the Fourth Pillar of Wikipedia policy, is the most useful indicator of what a Wikipedia policy actually is. It is formulated as follows:

If a rule prevents you from improving or maintaining Wikipedia, ignore it.

(Wikipedia: Ignore all rules)

This is a rather puzzling statement: if rules can be ignored, what are they there for? Ignore All Rules stipulates one can ignore rules as long as one is doing so in the interest of improving Wikipedia. What "improving" means is not defined: the judgment is left to the editor, who can decide what rules

to break and how. Of course, if their behaviour is considered wrong by the rest of the community, or by people in the Wikipedia hierarchy, the editor in question could be banned, or their edits reverted. It should be noted that not all rules can really be ignored: NPOV is unchangeable. A newcomer might not know this: what exactly is set in stone and what is up for debate can only be learnt by existing within the community, observing the behaviour of more experienced editors. Knowing when and how to break a rule – and consequently, how to follow the rule Ignore All Rules – appropriately is a skill that is learnt over time and with experience.

Wikipedia's policy should be understood in conjunction with the fact that "Wikilawyering", i.e. citing a policy in a discussion, wielding it against one's opponent (Wikipedia: Wikilawyering) is frowned upon. In sum: rules are not there to be respected, and neither can be appealed to when needed. In my view, Wikipedia policies are not policies: their role is to outline a worldview, signalling to newcomers what values are dear to members of the community. In doing so, policy encourages those who agree with those values to join the community, while discouraging those who don't. The frequent appeals to common sense and personal judgment – implicit in Ignore All Rules, as well – go in the same direction. Understanding what is and is not appropriate when it comes to both upholding and breaking a rule is essential to be an effective Wikipedia editor, and the kind of knowledge that is acquired with experience (Ford & Geiger, 2012). Hence Ignore All Rules, in practice, signals that, if you are able to navigate Wikipedia's social system, you can do whatever you want. Gaming the system, while frowned upon (Wikipedia: Gaming the System) is common practice: even getting out of being accused of gaming the system, by managing pushback from other editors, can be achieved by gaming the system. Wikipedia's policies, then, essentially encourage gaming the system.

Civility is the only conduct policy included in the Five Pillars. As such, I take it to be the most fundamental principle as far as regulation of behaviour is concerned. An outright definition of civility is not provided on the page – presumably, working on the assumption that it is not needed. The closest the policy gets to describing civility is the following:

Editors are expected to be reasonably cooperative, to refrain from making personal attacks, to work within the scope of policies, and to be responsive to good-faith questions. Try to treat your fellow editors as respected colleagues with whom you are working on an important project. Be especially welcoming and patient towards new users who contribute constructively, but politely discourage non-constructive newcomers.

(Wikipedia: Civility)

Other conduct policies and guidelines appear to be inspired by Civility, either as forbidding behaviours that are considered uncivil (e.g. harassment, listed in as a breach of civility) or as prescribing behaviours that are in line with conducting oneself in a "civil" way as outlined in "Wikipedia: Etiquette" (Wikipedia: Etiquette).

Underlying Civility and its corollaries is a culture founded on competition, or even regimented aggression. Wikipedia, as a system, doesn't discourage anger – the roots of this leniency being in the self-described anarchic environment created at its inception; initially, founders thought that policing behaviour would discourage participation (Sanger, 2005). In fact, while editors expressing feelings

of attachment, hope, or love in talk pages are harshly criticised, clearly frustrated – if not outwardly aggressive – behaviour is overlooked. The Civility policy condones anger – as long as it is not expressed through open aggression. In the section “Avoiding Incivility”, the Civility policy begrudgingly prescribes how to avoid “incivility – or the appearance of incivility” (Wikipedia: Civility). Some level of entitlement is condoned: “*Someone may very well be an idiot*, but telling them so is neither going to increase their intelligence nor improve your ability to communicate with them” (Wikipedia: Civility, emphasis mine). In the same section, it is also acknowledged that the Wikipedian assumed by the system – the user the policy is addressed to – is able to withstand aggressiveness. Those who can’t are marked as other.

Asking yourself ‘How would I feel if someone said that to me?’ is often not enough; many people can just brush things off. To get a better perspective, ask yourself: ‘How would I feel if someone said that to someone I love who *cannot* just ‘brush it off?’

(*Wikipedia: Civility, original emphasis*)

Throughout, the main concern doesn’t seem to be sincere respect for others. Rather, the Civility policy reads as a set of instructions for plausible deniability, detailing how to avoid being accused of incivility. Withstanding aggressiveness requires a lot of emotional labour (as theorised in Hochschild, 1983), and it comes easier to some than others. Again, this is a foundational quality of the project:

Wikipedia’s environment was not cooperative but instead competitive, and the competition often concerned what sort of community Wikipedia should be: radically anarchical and uncontrolled, or instead more singlemindedly devoted to building an encyclopedia. Sadly, few among those who would love to work on Wikipedia could thrive in such a protean environment.

(*Sanger, 2005*)

Additionally, Civility encodes conducts common amongst the professional class: Wikipedians are explicitly asked to behave as if they were at work. The policy uses a newsroom as an example; the standards for behaviour are lifted from office work (for a further exploration of this point, see Sundin, 2011). Since interactions regulated by Civility are mainly text-based, the use of language can make a difference between being considered civil or uncivil: native language, class and neurological makeup are among the factors that can influence how one is perceived. Hence, neurotypical native English speakers, with a certain level of education, able to withstand attack, and to act according to rules that are learnt outside of Wikipedia, and only in certain environments, have the upper hand.

In sum, Wikipedia policies outline a worldview, signalling to newcomers what values are dear to members of the community. In so doing, they encourage those who agree to join the community, while discouraging those who don’t. The frequent appeals to common sense and personal judgment – implicit in Ignore All Rules, as well – aim in the same direction. Personal judgment is not neutral. It is an expression of one’s point of view. At the same time, Wikipedia policy implicitly and subtly favours longstanding members of the community. Policy, then, functions as a device for subtly

selecting like-minded individuals, and Ignore all rules engenders a series of behaviours that construct and signal belonging.

1.3.3. Coded objects

In order to understand how Wikipedia works, it is important to understand how it's built. Wikipedia is a wiki: a website which can be edited through a browser. It is also a digital platform, powered by a system of programs, which are written and updated by a coding community, based on specific engineering principles.

Wikipedia's coded infrastructure can be thought of as being made of core software, surrounded by a cloud of small, but essential programs, known as bespoke code (Geiger, 2014). The core of Wikipedia's code is a piece of software called MediaWiki, which was developed specifically for Wikipedia, but is open source – hence it can be used for other wiki projects (MediaWiki: Manual: MediaWiki architecture). Bespoke code comprises a multitude of different programs, written in various languages, which perform a wide array of actions, some of which are necessary to the functioning of Wikipedia (Geiger, 2014). For instance, the “citation needed” tag works based on a script known as a template, stored separately from core code (*ibid.*).

The most salient pieces of bespoke code, for present purposes, are bots and gadgets. Bots are small programs that perform specific tasks, such as creating articles based on existing material, checking the website for expired links, and detecting potential vandalism. They can be very powerful, and as such are strictly regulated: bots have to be approved before they can be deployed (Wikipedia: Bot policy). Gadgets are written in JavaScript or CSS and typically used to improve the editing experience: they work like add-ons, that users can activate from their user preferences, and perform actions such as opening a link in a new tab, automating routine tasks such as welcoming new users, navigate quickly between page revisions, and so forth (Special: Gadgets). Some gadgets, however, can be made default, becoming, in practice, part of the standard experience of Wikipedia. Gadgets may seem mostly cosmetic, but can also interfere with key functionalities (I will describe an example of this in chapter 4).

It is important to note that Wikipedia's coded infrastructure is created by both professional engineers, employed by the Wikimedia Foundation, and volunteer programmers. In the case of bots and gadgets, anyone can create one, and submit it to the community for approval: in other words, volunteers create them of their own initiative, depending on their needs or desires. The core of Wikipedia's infrastructure, however, is developed based on a centrally established strategy, with input from the community, through a system of management called Phabricator. Phabricator is an online platform used to distribute programming tasks to whomever wishes to contribute. Volunteer developers with an account can pick up a task, work on it, then submit code for review on a platform called Gerrit (MediaWiki: Phabricator/Help). Advanced contributors can also set up tasks, subscribe to specific projects, suggest improvements. The Wikimedia Foundation also employs engineers to work on core code.

The values that inspire Wikipedia's development are laid out in an official Wikimedia development policy document, the “Wikimedia Engineering Architecture Principles” (MediaWiki: Wikimedia Engineering Architecture Principles). The document lists values and broad areas of application, indicating how they should be operationalised. The principles are (all quoted sections

from the policy document itself):

- Equity: “To allow users to consume, create, and interact in a form suitable for their devices, with the connectivity they have, in a language they speak”
- Empower: “To empower contributors to collaboratively grow and curate content, and to build the tools that they need to do so”
- API: “To provide public APIs that allow efficient interaction with wiki content, as well as provide data that can be easily processed and reused in bulk”
- Open: “To provide an open-source software stack that can be easily used, modified, and extended by others”
- Solid: “To maintain a code base that can be modified with confidence and readily understood”
- Run: “To provide a Web application that can be freely used to collaboratively collect and share knowledge”
- Fast: “To ensure availability and performance of WMF projects through scalable and resilient system design”
- Secure: “To ensure the data integrity of the content on WMF systems, and protect the privacy of our users”

As a set, the principles match the Wikipedian ethos of openness, universality and empowerment of individuals: software should be designed in such a way that allows participation of as a broad a number of people as possible, autonomously, and with as few constraints as possible. A mix of technical and political considerations inform these principles: being able to edit on any device, for instance, as well as being generally convenient, means opening up Wikipedia editing to people residing in countries where the Internet is mostly access through mobiles. More subtly, API and Open enable certain mechanisms of internal Wikipedia politics, as they are necessary for forking, i.e. downloading the whole of Wikipedia’s data, except user data, and start another site from scratch. Forking can be used as a form of dissent, as in the case of the of Spanish Wikipedia, which forked over a disagreement in 2002 (Reagle, 2010; Tkacz, 2011).

The document also details directions for operationalisation: the principle of Equity, for instance, applies specifically to devices, language and accessibility, rather than, say, content formats or technical skill. Operationalisations, to an extent, define the principles themselves: a broad value such as equity, when applied to design, can be interpreted differently depending on context. Deciding that equity applies to devices, language and accessibility, in the context of a collaborative encyclopaedia, implicitly defines equity as the ability to contribute regardless of geography or the specific kinds of disability that tend to be associated with online accessibility. The specificity of this notion of equity becomes apparent when one considers other variables such as gender, race, neurodivergence, or non-Western epistemologies, none of which is addressed by the definition of equity as operationalised in the Wikipedia engineering principles. I will expand on this point in chapter 5.

As well as being governed by policies, Wikipedia’s knowledge-production process is channelled by code: editing Wikipedia is made possible by a set of technologies developed collectively, and based on a specific set of values. Conversely, engineering values are co-defined by their

implementations, and Wikipedians are not just users of software created elsewhere: some of them are coders themselves. As I will articulate over the course of the thesis, the complex relationship between code, volunteers, and salaried leadership questions established assumptions the connection between values and design.

1.4. How Wikipedia works, in theory

Describing the apparatus that sustains the production of Wikipedia content doesn't fully explain Wikipedia's success, insofar as it is successful, nor its shortcomings. In this section, I will give an overview of existing epistemologies of Wikipedia, both resulting from reflexive work by editors themselves, and scholarly analysis. I will then circumscribe the gap in knowledge that the present study seeks to address.

1.4.1. Wikipedian epistemologies of Wikipedia

When Wikipedians discuss epistemic matters, they tend to do so mostly implicitly, by discussing practice, and through essays concerning specific aspects of their own work, but mostly refrain from explicit explanations of Wikipedia's functioning, with some exceptions. The few overarching, native narratives about Wikipedia's functioning adopt brute force as a core explanation. "Darwkinism" stipulates that the fittest information survives, after undergoing a form of natural selection (Meta: Darwkinism). Some Wikipedians appear, similarly, to envisage a sort of invisible hand governing the apparently chaotic activity of the editors, assuming that content will improve over time. In any case, Wikipedia is seen, epistemologically, as a competitive environment, in which the more accurate information wins. Wikipedians recognise themselves as part of a community, but hardly see cooperation as a way to produce knowledge – rather, they take competition to be the mechanism that sustains the continuous process of pruning of incorrect or misleading information.

"Eventualism" also highlights an important feature of knowledge production on Wikipedia: eventually all information will be corrected, so just leave it to the collective to take care of it. Faith in eventualism also implicitly justifies the extreme openness of Wikipedia: it may be vulnerable to vandalism, but it is assumed that it will recover from any attack, eventually. When "eventually" will come is never discussed. Given the constant developments in knowledge, this process is potentially infinite. In Wikipedia time, the difference between discussing a change before or after making it is minimal, as precision at any given time is not valued as much as some form of ultimate accuracy, achieved through collective work. Progress, here, is a symmetrical notion: it doesn't matter whether false information is included in the content of an article. Its elimination is in the hand of the editors, who will, at some point in the future, correct it. Equally, if true and relevant information is eliminated, it is assumed that, eventually, someone will notice and right the wrong. Were articles not so malleable as to allow easily for error, they wouldn't be open to accumulate correct information either. Precision, when achieved, retroactively justifies any past mistake, and is immediately in jeopardy again.

As pointed out in section 1.2, some accounts of Wikipedia are written by scholars who also belong to the community. For the most part, those accounts are descriptive, and while they might imply evaluations of Wikipedia, they rarely explicitly try to explain it. As far as I can tell, only one

author attempts to do so: Joseph Reagle, in *Good Faith Collaboration* (2010). Reagle (2010) attributes Wikipedia's success in creating a universal encyclopaedia to the way in which the Wikipedia community collaborates, and specifically to upholding two principles: neutral point of view and good faith collaboration.

I have already explained the neutral point of view policy. Good faith collaboration refers to the principle that Wikipedians should assume that each other's actions are done in good faith (Good faith). According to Reagle, adherence to these principles is imperfect in practice – they are to be treated like inspirational ideals. Regardless, insofar as they are present in the community, they allow for fruitful collaboration, which in turn sustains the production of knowledge on Wikipedia (*ibid.*). The mechanisms by which neutral point of view and good faith allow for effective collaboration are not thoroughly explored. Reagle writes that “Whereas NPOV renders the subject matter of a collaborative encyclopedia compatible, good faith makes it possible to work together” (p. 71). He seems to say, then, that the absence of a point of view erases any ideological reason not to contribute to the encyclopaedia, a line of reasoning that can be traced back to the founding of Wikipedia itself, and Sanger holds in his memoirs (Sanger, 2005). Good faith, on the other hand, together with other rules for civil behaviour, allows to reduce friction within Wikipedia (Reagle, 2010).

The problem with Reagle's account is that it doesn't concern itself with practice: it simply takes at face value the principles described in Wikipedia's policy, and claims that they achieve what they set out to achieve. Any behaviour within the community that contravenes to its principles is treated as an exception; after all, the principles are aspirational in nature. Reagle doesn't question the role of policies, ignores the implicit loophole offered by Ignore all Rules, and overall doesn't engage critically with Wikipedia's ideology at all. What's more, he barely engages with Wikipedia as a technological object, as other authors (section 1.4.3, below) have done: he mentions technology in passing, only to dispel the deterministic fantasy that using wiki technology would be sufficient to create a complex project such as Wikipedia. While he is of course right to say so, he is battling a straw man: no one, to the best of my knowledge, is attributing Wikipedia's success to its platform alone. On the other hand, downplaying the technological aspects of knowledge-production on Wikipedia to the point of mere facilitators of a social process is also, as others have argued (see section 1.4.3), naïve. In the next section, I will sketch how scholars who are not members of the Wikipedia community have tried to make sense of Wikipedia as an epistemic enterprise.

1.4.2. Independent, scholarly epistemologies of Wikipedia

In this section, I will give an overview of the account of Wikipedia's epistemology that have been published by academics. I will start with the more abstract treatments of philosophers, and work my way down to those authors that have engaged with Wikipedia's materiality and design.

The main concern of philosophers writing on Wikipedia has been the issue of trustworthiness: can we trust Wikipedia? The question of trust assumes two forms: it's a question of accuracy, concerned with whether Wikipedia's content is true; it's also a pragmatic question, concerning whether (or how) we should use Wikipedia as a source of information. All authors seem to agree that, since anyone can edit Wikipedia, expertise cannot constitute a foundation for its reliability: rather, the process of knowledge-production is the guarantor of truth, insofar as truth is achieved. The fact that Wikipedia is edited by a large group of people, for some authors, plays in its favour: whether it's

the wisdom of crowds and the Condorcet theorem (Fallis, 2008), or a mechanism similar to the invisible hand (Wray, 2009), or Darwinian competition (Smart, 2018), Wikipedia's strength is definitely found in numbers. Others (Fricker, 2011; Tollefsen, 2012) make a distinction between a crowd – a sum of individuals – and a community – a group of people bonded by personal relationships, purpose, and a sense of responsibility for their collective endeavour. If Wikipedia is a community, these authors argue, then it can be trusted to produce reliable information.

Epistemic virtues and vices also play a role in the accounts above. Where Wikipedia's success is pinned on editors constituting a community, virtues such as testimonial trustworthiness and accountability matter greatly (Fricker, 2011, Tollefsen, 2012). Accounts framing Wikipedia as a crowd, on the other hand, either ignore virtues altogether (Fallis, 2008; Wray, 2009) or leverage vice as a means of getting unstuck from local maxima (Smart, 2018). Having established, one way or another, whether Wikipedia is epistemically sound, then informs how Wikipedia can be used, with some authors arguing for moderation, using Wikipedia as a starting point for further research (Magnus, 2009), as entertainment or to find community (Fallis, 2008).

Philosophical treatments of Wikipedia tend to engage with epistemic values such as accuracy and trustworthiness, but neglect to ask questions around epistemic justice, and ignore the technologies that mediate editing. Others (e.g. Antin *et al.*, 2012; Ford, 2022; Friedenthal, 2017; Geiger, 2014; Kukowska, 2022; Niederer & Van Dijck, 2010) have rightfully pointed out that Wikipedia's functioning is made possible by a complex interaction between the community, technologies and procedures: that Wikipedia is, in other words, a sociotechnical system.

The most extensive body of work on Wikipedia as a sociotechnical system has been carried out by Heather Ford, who has published various works on the subject. Ford connects social epistemology and platform affordances, i.e. the actions the platform lends itself to (see Gibson, 1979, for the original definition of affordance). Ford conceptualises Wikipedia's knowledge-production process as made possible by an apparatus of technologies and people that shepherd facts from external sources to Wikipedia, and from there onwards, to other virtual spaces where those facts are re-used and repackaged, such as search engines (Ford, 2022). At the same time, the travelling of facts obscures their provenance, thereby hiding potential sources of bias (Ford, 2022).

Unsurprisingly, since materiality and epistemic justice are long-term allies, analysing the ways in which Wikipedia's infrastructure regulates knowledge-production has also brought to the fore issues around justice. It has been argued that technical skill is a component of Wikipedia's specific notion of expertise, as editing content requires an understanding of how the software works (Ford & Geiger, 2012; Hartelius, 2010) and the ability to code is at least beneficial, as it allows editors to expand their influence (Geiger, 2014). The sociotechnical affordances of the Wikipedia editing platform filter out people with little technical skill, which often happen to be women (Ford & Wajcman, 2017), and grant more influence to those who are already part of the community (Ford & Geiger, 2012).

The design of Wikipedia's coded infrastructure is in and of itself meaningful. Maja van der Velden has pointed out how the emphasis on propositional language, and the categorisation system on Wikipedia makes it inhospitable to indigenous knowledge (Van der Velden, 2013). Others have studied specific coded components of Wikipedia, in terms of their impact on community life (Morgan & Halfaker, 2018), their behaviour (Tsvetkova *et al.*, 2017), as embedding moral values (Geiger &

Ribes 2014; Kuznetsov, 2006).

Geiger (2014) points out how taking the materiality of Wikipedia’s infrastructure seriously essentially means questioning how power and code relate to each other. The dictum “Code is law”, popularised by Lawrence Lessig and widely used in theoretical work about platforms, doesn’t fit Wikipedia (*ibid.*). In Lessig’s account, code works as an abstract structure that imposes certain kinds of behaviour on its users; this, points out Geiger, obscures the material process of production of code, which, as I have described above, is rather complex, and gives an illusion of harmony, of code as a unifying principle that shapes action. By foregrounding the process of coding, Geiger expands the field of analysis: affordances – how coded objects are used, and allow to be used – are only part of the picture. The process of coding is interesting to Geiger, as far as I can tell, mostly in terms of power: being able to code is, in his view, necessary to edit Wikipedia. I agree, and will leverage this insight over the course of my analysis. I will, however, push further, and look at the material agency of code itself, to the consequences of the way it is structured, written, organised.

In terms of engagement with materiality, studies of Wikipedia’s epistemology sit on a spectrum, from those who almost completely ignore it (Fallis, 2008; Wray, 2009), to those who focus on the relationship between the social and the technological, specifically in terms of use (Antin *et al.*, 2012; Ford, 2022; Niederer & Van Dijck, 2010), to those who focus primarily on the material, as an extension of the social (Geiger, 2014). My study, as I will outline in chapter 2, stretches said spectrum further towards a material analysis: while of course I will acknowledge, and explore, the social aspects of designing and maintaining Wikipedia, I will aim to take seriously the idiosyncratic form of agency of Wikipedia’s infrastructure, which leads it to have effects beyond merely implementing values held by its designers. Taking materiality seriously entails engaging with the specific design affordances of digital technologies, and how they intersect with the practice of coding. This is particularly important in this case, because, as I will show in the next section, Wikipedia is a product of programming culture: founded and developed by people who were part of technological communities, inspired by ideals cultivated within the free software movement, and designed based, in part, on the philosophy of Ayn Rand, a staple Silicon Valley author. The ramifications of Wikipedia’s roots in tech culture, as I will argue over the course of the thesis, infiltrate the whole project, shaping its epistemology and design.

1.5. Wikipedia’s roots

In this section, I will trace an incomplete history of Wikipedia and its influences. Specifically, I will focus on hacker culture (and surrounding areas), a direct ancestor of Wikipedia; the philosophy of Ayn Rand, appreciated by some of Wikipedia’s founders; and encyclopaedias, for obvious reasons. I have left this historical sketch for last, because I want to be able to draw attention to the connections between historical influences and Wikipedia’s culture and characteristics, described above. These connections lay the groundwork for understanding how and why Wikipedia has taken its current form.

1.5.1. Historical overview

Wikipedia was founded in 2001, either by Larry Sanger and Jimmy Wales, or Jimmy Wales,

depending on sources; Wikipedia started as an auxiliary project of Nupedia, a collaborative online encyclopaedia staffed with experts (Sanger 2005, Lih 2009). Initially, non-experts involved via Wikipedia were supposed to support Nupedia's experts, but Wikipedia turned out to be more efficient at creating articles, due to the absence of peer-review, thus winning out in the end (Sanger, 2005).

When Nupedia was developed, experimentation with collective knowledge sharing on the Web was already underway, with projects such as Interpedia and Project Gutenberg (Lih, 2009). More specifically, the idea of a collaborative encyclopaedia was already in the air, with h2g2, created by sci-fi author Douglas Adams, and GNUpedia, a planned distributed encyclopaedia composed of websites and resources scattered around the Internet and linked together in a loose network, which was launched in 1999, and then abandoned when Nupedia was created (Free Software Foundation, 2019; Lih, 2009). The specifics of the relationship between Wikipedia and the free software movement will be discussed in section 1.5.3.

Nupedia, and then Wikipedia, were funded with the revenue from Bomis, Jimmy Wales's pre-existing company, a Web portal that hosted entertainment and adult content (Lih, 2009). Before founding Bomis, Wales had already worked in finance for a few years (Jemielniak, 2014; Reagle, 2010; Van Dijck, 2013). Jimmy Wales and Larry Sanger had met through an Ayn Rand discussion group (Reagle, 2010). When he was hired to develop Wikipedia, Larry Sanger had just finished a doctorate in philosophy. Sanger eventually left, in 2002, due to disagreements about the project (Sanger, 2005). He has been since a very vocal critic of Wikipedia, and even started a rival project based on different principles: Contropedia.

Wikipedia's overall development, in its first 20 years of existence, can be described as a process of hardening. Starting as a relatively supple project, open in terms of content, ability to contribute, conduct and technical development, Wikipedia's structure became increasingly rigid: categories ended up structuring its content, protections were established that put restrictions on what pages could be edited by whom, policies and guidelines were written, and the coded infrastructure became more and more fixed.

At the beginning, Larry Sanger didn't, as a matter of principle, try and direct the actions of early Wikipedians, who were given latitude to contribute as they saw fit. The process of hardening occurred somewhat haphazardly, in response to specific needs, and as a result of individual choices and interventions. Protections and semi-protections, for instance, which prevent a large part of the contributor pool to edit protected or semi-protected pages, were introduced in 2005, in response to an event known as the Seigenthaler incident, which entailed the vandalism of a living person's biography (Frost-Arnold, 2018; Reagle, 2010). Consequently, any structuring of processes, policies and so forth occurred without a central strategy or authority, ending up with a messy, yet hard to change, general setup.

Over time, Wikipedia became increasingly bureaucratic as well: the corpus of policies and guidelines grew, and an organisation responsible for Wikipedia's publication was established, the Wikimedia Foundation (WMF) (Fuster Morell, 2011). The Wikimedia Foundation, as well as hosting Wikipedia and providing other services such as fundraising, supports initiatives aimed at the free distribution of knowledge (Ayers *et al.* 2008). While the WMF employs engineers that work on Wikipedia's code alongside volunteers, including a Principal Software Engineer whose work is to

oversee development, it doesn't intervene directly in editorial matters, which are to this day left to volunteers. The WMF does, however, set strategic priorities for the project. Because of its focus on infrastructure and funding, said priorities tend to concern design choices, such as the format in which information is presented, and providing resources for events or outreach. The current strategy, "Wikimedia 2030", was established in 2021 (MetaWiki: Movement Strategy).

Similarly to content and policy, Wikipedia's coded infrastructure has been provided with new features, while becoming more and more rigid over time. In terms of the editing interface, the default interface was improved with a Visual Editor, a WYSIWYG interface that allows to edit pages without modifying markup (details and analysis in chapter 4), a smartphone editor, and integration with machine learning. The core of Wikipedia's code, having been developed without a pre-established plan, grew in a way that can be described as entangled: rather than being made of parts that fit a clear structure, such as, say, a piece of furniture, it was accumulated over time by aggregating various kinds of material, made by different people, as needed, resulting in something more similar to a block of granite – while technically a rock, granite is made of small bits of various minerals, all stuck together: solid, but heterogeneous, and lacking internal structure. This kind of coded structure is known, fittingly, as a monolith. Monoliths are very difficult to change, since the absence of a map of the internal relations between parts makes it difficult to predict the result of any intervention. I will expand on the consequences of this fact in chapters 3 and 4.

1.5.2. Hacker culture and the free software movement

Wikipedia has historical and ideological ties to the free software movement (or free and open software movement, from which the standard abbreviation, FOSS, is generated). The free software movement was started in 1983, 1984 or 1985, depending on accounts, by then-MIT programmer Richard Stallman. The movement is concerned with promoting the use and development of non-proprietary software. Stallman's first practical act was to create GNU, an operating system consisting entirely of free software. According to Stallman, software is "Free as in 'freedom', not 'free' as in 'beer'":

When we call software "free," we mean that it respects the users' essential freedoms: the freedom to run it, to study and change it, and to redistribute copies with or without changes. This is a matter of freedom, not price, so think of "free speech," not "free beer."

(Free Software Foundation, 2024)

Free software is free in that it can circulate freely, be adapted and reused by anyone, on the condition that any product created by using free software is free in turn, as in reusable, and so forth. Free software is also free in the sense of gratis.

Richard Stallman's work was an ideological ancestor of the Wikipedia community, and at some point, even a Wikipedia competitor, as acknowledged by on Wikipedia itself (Wikipedia: History of Wikipedia). Wikipedia uses free software, which, in and of itself, only indicates a broad allegiance to the movement, not an ideological link. A better starting point to look at Wikipedia as an heir of the free software movement is on Wikipedia's website: Wikipedia's tagline is "The free encyclopedia

that anyone can edit”. The word “free” in the tagline displayed on Wikipedia’s main page links to the page (Free content), which refers to the cultural version of free software, as defined by Stallman: content that can be accessed, studied, applied, copied, modified, and redistributed (Free Software Foundation, 2024).

Stallman, on his part, had his own encyclopaedic interests. Throughout his career, he has campaigned for open knowledge as well as free software. In this context, he sketched an idea for a free encyclopaedia, around the time Wikipedia was invented, centred around the idea of freedom. As mentioned above, his blueprint was about to turn into a real project, GNUpedia, when Nupedia adopted the GNU license; hence GNUpedia was abandoned (Stallman, 2024). Stallman considers Wikipedia the spiritual heir of GNUpedia, and the only free encyclopaedia.

The principles laid out by Stallman closely resemble some of Wikipedia’s values. In his plan for a free encyclopaedia, Stallman set out “criteria of freeness” (Free Software Foundation, 2024). The first one is a throwback to “free as in beer”. While for software, Stallman almost dismissed the monetary aspect, emphasising the ideological aspect, for the encyclopaedia he thinks specifically about profit:

Conventional nonfree encyclopedias published by companies such as Microsoft will surely be made available on the web, sooner or later—but you will probably have to pay to read an article, and you surely won’t be allowed to redistribute them. If we are content with knowledge as a commodity, accessible only through a computerized bureaucracy, we can simply let companies provide it.

But if we want to keep human knowledge open and freely available to humanity, we have to do the work to make it available that way. We have to write a free encyclopedia—so we must first determine the proper interpretation of “free” for an encyclopedia on the Internet. We must decide what criteria of freedom a free encyclopedia and a free learning resource should meet.

(Free Software Foundation, 2024)

He emphasises further that access has to be universal, at least for those who possess a computer. The second criterion for “freeness” is the ability to replicate the site:

Permit mirror sites.

When information is available on the Web only at one site, its availability is vulnerable. A local problem—a computer crash, an earthquake or flood, a budget cut, a change in policy of the school administration—could cut off access for everyone forever. To guard against loss of the encyclopedia’s material, we should make sure that every piece of the encyclopedia is available from many sites on the Internet, and that new copies can be put up if some disappear.

(Free Software Foundation, 2024)

The ability to duplicate information is, here, framed as a safety issue – not, as with software, as a matter of principle. While reusing and redistributing information makes ideological sense in a

framework that values access to knowledge, duplicating an exiting, available resource doesn't: access is granted to any number of people without a need to reproduce it. Stallman then carries on: the encyclopaedia should permit translation in languages other than English, permit quotations, and permit to modify and redistribute sections of the encyclopaedia. Evoking them here serves to illustrate the relationship between freedom and access, and to aid in reading the way in which the Wikipedia community conceptualises freedom as applied to knowledge-production. One last note on Stallman's treatment of the concept of freedom: throughout his writings, he keeps going back to the implicit question of governance. Every now and again, he answers a question no one has answered, and namely: is an apparatus (or, in his terms 'a bureaucracy') needed in order to achieve what has just been outlined? For instance, when talking about mirroring:

There is no need to set up an organization or a bureaucracy to do this, because Internet users like to set up "mirror sites" which hold duplicate copies of interesting Web pages. What we must do in advance is ensure that this is legally permitted.

(Free Software Foundation, 2024)

Wikipedia, in sum, has inherited, possibly via Stallman's outline, or perhaps simply by building on similar principles, large parts of the free software movement ethos. Wikipedia is free in the sense that it is gratis, i.e. people can access Wikipedia without payment; it is also free to move around, be used, changed, and redistributed. Finally, it is free in a meaning that aligns more closely with "free" in "free speech": Wikipedia is not censored (Wikipedia: What Wikipedia is not) and any topic, in principle, is suitable for inclusion.

In my view, Wikipedia's links to the free software movement also place it in the same cultural environment: tech communities. Wikipedia's founders, first and foremost Jimmy Wales, were active in technological circles. As mentioned above, Wales himself was an early Internet entrepreneur. Word of mouth about Wikipedia spread, initially, through online tech publications, which attracted a readership of amateur programmers and tinkerers, or, in other words: hackers.

I believe Wikipedia should be understood in the wider context of hacker culture, which developed from the 1970s onwards among computer programmers. The term hacker was initially used in the sense of a crafty person who is able to concoct clever systems in order to make machinery work (Levy, 1994). Over time, it ended up referring to people who, for profit or fun, infiltrated secure IT systems (Steinmetz, 2015). Alongside the hacker, the figures of the geek and the nerd became established as identities and cultural tropes: all three can be, and often are, applied to the same person, and have contiguous, if slightly different, meanings. "Geek" and "nerd" tend to refer to people who know a lot about something – which could be technology, but also any other topic (Kendall, 1999b; Steinmetz, 2015). Both also tend to refer to broader cultural traits, such as enjoying certain hobbies and genres of fiction, and personal characteristics, such as unattractiveness and strangeness: the word "geek" comes from a word that meant "carnival freak" (Levy, 1994). In cultural terms, the three words can be used mostly interchangeably, unless specific practices, such as picking locks, are concerned. Consequently, I will use them as synonyms, unless the context requires precision.

Hackers, nerds and geeks have a specific with knowledge. First, they all value it; knowledge is

precious currency in hacker culture. Secondly, specifically among hackers, there is a desire for knowledge and tools to be freely accessible; this has been true since early hacker culture, when coders would learn how to pick locks in order to access computer rooms out of hours (Coleman, 2013). Over time, the idea that knowledge should be free became attached to the practice of accessing databases. The free software movement grew intertwined with hacker culture, just focusing on the idea that both tools and information should be freely accessible. The third aspect of hacker culture that relates to knowledge is how it is spread: informally. Because hacking has to do with clever solutions to practical problems, it cannot easily be taught in an institutional setting; typically, it is learnt through peer interaction (Sennett, 2008; Steinmetz, 2015). In fact, the practice of teaching each other is one of the bases for the social aspect of nerd culture: forums became popular, in part, as ways of sharing and cataloguing programming tricks and tips based on experience (Steinmetz, 2015). The tight coupling of nerd culture and knowledge-gathering is some of the other encyclopaedic projects launched around the same time as GNUpedia and Wikipedia, in areas belonging to nerd culture, such as Memory Alpha, dedicated to Star Trek fandom, and the aforementioned h2g2, created by science fiction writer Douglas Adams.

The most obvious connection between hacker culture and Wikipedia is the ideal of universal availability of knowledge; detailing this further would be redundant. A point that requires more explanation is the emphasis on informal sharing of knowledge, and its counterpart, mistrust for institutions and authority. Wikipedia culture is averse to authority, privileging a decentralised system, and, as I explained above, downplaying the significance of power differentials even when in the presence of clear hierarchies. Consistently, given a context where knowledge is associated with privilege, subject expertise doesn't improve the chances of success of individual contributors: anyone can edit, and leveraging one's own real-world credentials is frowned upon (Appendix A). Conversely, knowledge acquired through peer learning is fundamental to thrive as an editor, as it allows to navigate, and take advantage of, the complex system of policies, interpretations, and unspoken rules which regulate editing.

Beyond epistemic commitments, hackers and Wikipedia share cultural traits more broadly. Both are characterised by a culture of competition geared around the product being created: for instance, in the case of programmers, coding challenges have been common since the 1970s (Ensmenger, 2015). Hacker culture is a masculinist, even misogynistic culture (Biricik & Hearn, 2009; Ensmenger, 2015; Ottemo, 2019); Wikipedia presents a similarly masculinist environment (Peake, 2015), and has been criticised for faring poorly in terms of gender equity as well (Ford & Wajcman, 2017). I have already mentioned anti-authoritarianism, above. The specific blend of commitment to one's craft, anti-authoritarianism, competitiveness, demographic uniformity, peer learning, and decentralisation found both in hacker and Wikipedia culture have exercised, in my view, significant pressure on Wikipedia's design. I will provide a full analysis in chapter 4.

Objects designed by nerd communities also tend to conform to a recognisable aesthetic, articulated in various forms sharing a performative disdain for traditional parameters of beauty (see for instance Turkle, 1984; Simon, 2007). I will argue in chapter 3 that Wikipedia's aesthetic, in terms of the design of its website, and the style of its content, is significantly indebted to nerd aesthetic.

Wikipedia, in conclusion, can be framed as a product of the free software movement and of hacker culture more broadly, regardless of Richard Stallman's specific role in its creation. In the next

section, I will describe how Wikipedia shares in another widely interest of tech communities: the work of Ayn Rand.

1.5.3. Ayn Rand

Wikipedia's ethos is rooted in the work of writer and philosopher Ayn Rand (1905–1987). Wikipedia's founders, Jimmy Wales and Larry Sanger, had met through a mailing list started by Wales called "MDOP: 'Moderated Discussion of Objectivist Philosophy'" (Schwartz, 2015), in which the works of Ayn Rand were discussed. Jimmy Wales got Larry Sanger involved because he wanted a philosopher to lead the project: "They saw the Nupedia project as turning objectivist theory into practice" (Lih 2009, p. 37). Sanger and Wales met through a mailing list, moderated by Wales, called "Moderated Discussion of Objectivist Philosophy" (Jimmy Wales; Schwartz, 2015). In his memoir, Larry Sanger notes that there was nothing inherently inevitable in the development of Wikipedia; rather, he and his collaborators made choices based on their own values (Sanger, 2005). Those values happened to align with the basic tenets of Objectivism, which they, hence, adopted as a philosophical guide.

Ayn Rand rarely appears in mainstream scholarship. Her work is used and taught mainly by a handful of specialised institutes, whose activity can be classed as apologetics. One of them, the The Ayn Rand Institute, is run by Leonard Peikoff, one of her former students. So was *The Objectivist*, a journal dedicated to the dissemination of Ayn Rand's ideas. *The Journal of Ayn Rand studies* appears to be the only outlet that publishes a variety of perspectives on her work.

Because of her relative academic obscurity, it seems fitting to briefly introduce her, and her work. Ayn Rand was born Alisa Zinovievna Rosenbaum, in St Petersburg, Russia, in 1905, from a relatively privileged family (Gladstein, 2010). She moved to New York in 1926, never to return again; her experience of life under the nascent Soviet Russia inspired in her a strong disdain for collectivism (*ibid.*). Once in the United States, she slowly and unsteadily built a career as a writer for herself, working as a screenwriter and publishing novels. Anti-collectivism is a recurring theme through her work, starting with *We the living*, written in order to show Americans the shortcomings of communism, during the so-called "Red Decade", a time when American intellectuals showed interest in, and commitment to, left-leaning ideals (Gladstein, 2010). The tenets of her political views were individualism and "laissez-faire capitalism" (Rand, 1943). When she died, in 1979, "a big floral arrangement in the shape of a six-foot tall dollar sign graced her funeral" (Gladstein, 2010, p. 20).

Objectivist epistemology is based on the idea that the world is a collection of entities, and that knowledge is a process of absorbing facts from the world directly, without any active participation (Rand, 1967/1990). According to Rand, facts are found, not constructed: "'fact' is simply a way of saying 'this is something that exists in reality' – as distinguished from imagination or misconception or error" (Rand, 1967/1990, p. 241). This definition doesn't leave space for points of view: it assumes that such a univocal apprehension of reality is possible in the first place. The neutral point of view can be, then, reasonably connected to the founders' engagement with Rand.

Ayn Rand's epistemology is based on her metaphysics. She claims that the world is made of existents, which are the unities of reality, and that facts are known existents (Rand, 1967/1990). Self-evident facts are "axiomatic concepts" or "axioms" (Rand, 1967/1990, p. 55). Axioms are the basis of objectivity, as knowledge is built upon them. In polemic with her contemporaries, which at

the time she was writing (1969–71) included existentialists and logical positivists, she contends that concepts of reality are derived directly from reality itself, and apprehended through reason. Knowing, according to Rand, entails defining concepts and putting them in relation to each other, creating a model of reality. Knowledge is the cumulative sum of all facts.

It has been noted elsewhere that Wikipedia is based on the idea that objective knowledge can be obtained through measurement and perception (Lih, 2009, Westerman, 2009). Facts are out there, waiting to be picked and collected into a comprehensive whole. The idea that an encyclopaedia should be curated and have a perspective is meaningless, within this framework. Editors are not creating a view of reality; rather, their job consists in representing reality as faithfully as possible, in the only possible way. The idea of different perspectives on reality is present in the form of controversies, but these are framed as disagreements that can eventually be settled, as there is a matter of fact, out there, waiting to be discovered.

This thinking is reflected in Jimmy Wales's Foreword to Lih's *The Wikipedia revolution* (2009). The Foreword opens with the following quote, from an interview given by Wales himself: "Imagine a world in which every single person is given free access to the sum of all human knowledge. That's what we're doing" (p. XV). This quote also summarises Wikimedia's vision, and is displayed on their website. How Wikipedia works, written by three Wikipedia editors as a user manual of sorts, opens with a quote from World Brain, by H.G. Wells: "A complete planetary memory for all mankind" (Ayers, Matthews, & Yates, 2008). The basic idea here is that knowledge is cumulative, and can be pooled together from available sources, in an additive process. Error is removed through subtraction, by reverting or deleting edits.

The reception of Ayn Rand's work has been extremely heterogeneous. Literary critics have mostly ignored her at first; *The Fountainhead* gave her some exposure; after the publication of *Atlas Shrugged* (1943), it became impossible to ignore her. She developed a solid group of acolytes, who eventually started developing her philosophy, Objectivism. Libertarians often credit her as an inspiration; she tended to disavow her relationship to libertarianism and neoliberalism, claiming that they had appropriated her philosophy (Gladstein, 2010). Her relationship with the feminist movement was of mutual distaste (Gladstein, 2010). However, parts of 60s and 70s American counterculture saw her as a kindred spirit, despite her dislike of them, perhaps because they encountered her through her novels (Daub, 2020). If the hard theoretical core is ignored, it is easy to see a family resemblance between hippy ideals and Ayn Rand's general attitude towards self-realisation, and against elites.

Counterculture is the channel that brought Ayn Rand to modern tech culture, and in turn, to Wikipedia. The ties between contemporary Silicon Valley and counterculture are strong and well-documented (Daub, 2020). Some parts of her corpus resonate with the go-getting culture of Silicon Valley: the idea of the creative entrepreneur leading social change (*Atlas Shrugged*), that school holds back innovation (*The Fountainhead* prefigures the trope of the Stanford dropout: its central hero, Howard Roark, is kicked out of architecture school for his bold ideas), the faith in unregulated markets, the proud dismissal of government and regulation. Despite the free software movement disavowal of big corporations, there is a strong kinship there – the cultural entanglements between Silicon Valley and the free software movement are manifold (Daub, 2020).

Ayn Rand's philosophy is explicitly individualistic. Her ethics is rooted in self-interest. Self-

interest is aestheticized in the form of a modern version of the romantic hero. In *The Fountainhead*, for example, the main character, Howard Roark, is a defiant young architect who fights against the traditionalist architectural establishment, for the affirmation of his style – a unique, personal style, that doesn't blend with the buildings of the past, like those of his predecessors. Roark is a romantic hero. As Daub (2002) noted, Ayn Rand's philosophy is centred around an aesthetic of genius: even the eminently collective activity of creating a building is reframed as the achievement of a single, exceptional human being. In this sense, Ayn Rand's choice of architecture as the setting for her novel is appropriate: at the time when she was writing *The Fountainhead*, architects were celebrated as artists, for their own, unique, vision.

While authorship on Wikipedia is firmly collective, there are traces of this way of thinking in its policies, and, consequentially, the way Wikipedia's content is produced. The policy "Be bold" similarly exalts the activity of the single editor, as a hero: someone who doesn't ask for permission to realise their vision. Wikipedia is a do-ocracy: editors don't wait for directives before acting; they act, and acting legitimises their contribution. "Ignore all rules" acts in the same vein, by giving primacy to an ethical imperative (the creation of content for Wikipedia, which is framed as a mission with strong ethical connotations) over bureaucracy. Generally speaking, Wikipedia's ethos is electively – certainly not practically, given the large number of rules, checks and balances – anti-bureaucratic, in line with its Randian, and countercultural roots.

1.5.4. Encyclopaedias

Wikipedia is an encyclopaedia. It inherits features, aesthetics and practices from its predecessors, partly through cultural influence – a certain idea of what an encyclopaedia should be and look like – partly as a result of the constraints imposed by the format – for example, the content of modern encyclopaedias is organised in a way that facilitates consultation – and finally, through the options that are allowed or precluded by the technological setup of Wikipedia. At the same time, being an encyclopaedia doesn't entirely define Wikipedia: the way Wikipedia is designed is a result of the intermingling of influences (see above) with the constraints of being an encyclopaedia, and the choices of its creators.

The aim of this section is not to squeeze a full history of encyclopaedias into a few paragraphs; rather, to sketch, in broad strokes, the range of options available when designing an encyclopaedia, and the method by which encyclopaedias are typically written, thus placing the choices made by the founders of Wikipedia in a wider context.

While it seems impossible to track down the oldest encyclopaedia (sources disagree and are mostly rather defeatist on this point), it is clear that encyclopaedias – or something like encyclopaedias – have existed for a very long time. There is evidence of a compendium of knowledge from the second millennium BCE, entitled *A text to dispel ignorance about everything that exists*, by an Egyptian author named Amenemope (Jackson, 1977); Pliny the Elder wrote his *Natural History* in the first century CE (Jackson, 1977); the *Huanglan* (which means "Imperial Antology") was written in China around 220 CE (Bauer, 1965).

In the historiography of encyclopaedias, modern encyclopaedias are characterised in terms of an objective attitude towards the matter at hand (König & Wolf, 2013), starting with the anonymous *Compendium Philosophie*, dated circa 1300 (Collison, 1964), and ambitions to completeness, which

became apparent in the production of encyclopaedias from the 1600s onwards, Diderot and D'Alembert's *Encyclopédie* being the most famous early example (Holmberg, 2021).

The most important technological innovations, from the point of view of making encyclopaedias, have been the printing press, which has fundamentally expanded the audience of encyclopaedias, and consequently changed the way they were written (Blair, 2006); the introduction of indexes, pioneered by bishop Antonio Zara in his *Anatomia ingeniorum et scientiarum*, in modern-day Croatia (Collison, 1964), the progressive adoption of the alphabetical order, become widespread in the 16th and 17th century (König & Woolf, 2013) and finally digitalisation, on physical supports such as CD-ROM and DVDs, and online (*ibid.*).

To these innovations, I would add the invention of the wiki – that is, a website that can be edited through a Web browser. The wiki has made possible a radical shift in the practice of making encyclopaedias, from increased access to consuming knowledge to increased access to producing knowledge. If until the late 20th century the readership of encyclopaedias has expanded, the 21st century has seen a manifold increase of the authorship of encyclopaedias, pioneered by Wikipedia.

Initially, encyclopaedias were supposed to be read in their entirety: they were conceived as educational courses; reading them was supposed to be a way to improve oneself, to become a fit citizen, a proper adult. One such example is Marcus Terentio Varro's *Libri IX disciplinae*, probably modelled on the traditional framework of trivium and quadrivium (Jackson, 1977). Stallman's sketch for an encyclopaedia, described above, is based on that template. Eventually, with the introduction of devices that made searching an encyclopaedia easier, they became reference works: readers were not expected to read them in their entirety – rather, to consult them as needed. A more active role was granted to the reader, who was also encouraged to find their own path through the facts presented in the encyclopaedia. This was made possible by cross-referencing devices such as 'see also', or arrows, available from the 15th century onwards (Collison, 1964).

Conceiving of encyclopaedias as reference works, as opposed to courses, influences their size as well: concision is not a necessity, but is sometimes a desideratum. In Wikipedia's case, the debate over Wikipedia's size – and consequently, the breadth of knowledge it covers, and the discussion over criteria of introduction of specific entries (notability) – is still live, pitting inclusionists, who advocate for a laxer attitude towards the inclusion on Wikipedia of topics that are deemed less important, against exclusionists, who support a more selective approach. This debate has consequences in terms of justice: it has a bearing on the coverage of people or topics that are underrepresented in the Western, male-dominated source literature. It also influences design choices, such as article length.

While a recurring narrative about encyclopaedias emphasises their objectivity – or at least, ambition to objectivity – the history of encyclopaedias shows a variegated landscape. Early encyclopaedias had a tendency towards a specific perspective on reality; Pliny is a vocal author, filling his pages with commentary (Jackson, 1977). Even though often the passage to modern encyclopaedias is framed as marked by attention to objectivity, there are notable exceptions. The *Encyclopédie* is politically vocal (*ibid.*). Novalis's unrealised encyclopaedic project, *Das Allgemeine Brouillon*, should have been, programmatically, a subjective encyclopaedia, from his point of view, mixing poetry and science (Novalis & Wood, 2007). While the genre of the modern encyclopaedia is mostly, and for good reason, presented as an offshoot of the enlightenment – and indeed Wikipedia

upholds the same ideals of rationality and primacy of science that inspired enlightenment encyclopaedias – other ways of making encyclopaedias have always existed, and still do. There are even interesting cases in which partiality and impartiality are mixed: the *Enciclopedia Italiana*, written during Fascist rule, was written in a relatively objective way, despite the direct involvement of Mussolini himself, who wrote the article on fascism (Turi, 1979). A modern example of an explicitly ideological encyclopaedia is Conservapedia, which describes itself as a “conservative, family-friendly Wiki encyclopedia”, created to counteract Wikipedia’s perceived anti-Christian, liberal bias (<https://www.conservapedia.com/Conservapedia>). Wikipedia’s emphasis on neutral point of view, then, is a choice – not a simple inheritance from the genre.

Encyclopaedias have more in common with libraries than novels. They are, by their own nature, tertiary literature, concerned with reporting what has been discovered elsewhere. In antiquity and up to the Middle Ages, they made heavy use of authority (Collison, 1964). Often, encyclopaedias have been re-workings or translations of earlier encyclopaedias (Holmberg, 2021). A certain self-effacement is visible throughout: with the aforementioned exceptions, authorship tends to take a backseat. Encyclopaedias are also collective endeavours: if the ambition is to cover as much of human knowledge as possible, then a person alone is unlikely to succeed. Writes Francis Bacon in the *Proemium to the Instauration Magna*:

he resolved to publish at once so much as he has been able to complete. The cause of which haste was not ambition for himself, but solicitude for the work; that in case of his death there might remain some outline and project of that which he had conceived, and some evidence likewise of his honest mind and inclination towards the benefit of the human race.

(Bacon, 1620/1909 – 14)

Unfinishedness is now accepted as an inherent feature of encyclopaedias: first there was the practice of issuing supplements, now of regularly updating websites. The epistemic attitude of encyclopaedias has shifted, from framing knowledge as something that was attained in the past, to a future, likely unattainable, goal (Holmberg, 2021).

Encyclopaedias have ambitions to universality and comprehensiveness (König & Woolf, 2013). The spatial and temporal vastness of encyclopaedias, together with their character of tertiary literature, corresponds to a humble conception of authorship, where merit is not given to the individual, but to a whole system. Up until the late 20th century, it was not customary for authors to sign encyclopaedia articles (Collison, 1964). Of course this is not true of all encyclopaedias and all articles – and the exceptions are telling, important and frequent – but such a widespread practice of anonymous writing is a rather unique feature of the encyclopaedic genre. Its significance becomes even more apparent when one considers the culture of authorship that is pervasive in academia, where one’s name attached to a piece of writing is vital, and originality essential in order to occupy a space. Wikipedia’s ethos around authorship is in keeping with this tradition of modesty: as described above, Wikipedia operates under a regime of radical collaboration, where no one’s contribution is untouchable, or more important than others, and the power of those with more access than others is often downplayed.

Similarly, encyclopaedias are referred to by the names of their publishers (Microsoft Encarta) or

editors (the *Encyclopédie* is known as Diderot's encyclopaedia). Crediting editors is not unwarranted, as it might be argued that, for project as open-ended, often programmatically impartial, and reliant on previous literature as modern encyclopaedias, the setup created by the editor(s) is as close to traditional authorship as it gets. The editors decide the scope, ambition and perspective of the encyclopaedia; who gets to contribute; what is knowledge and what isn't; how knowledge should be organised; how contributors should collaborate. Myriad value judgments and decisions are involved in this process, and the responsibility for those judgments' rests on the curators: when controversy erupted around the content of the *Encyclopédie*, Diderot alone was imprisoned – not other authors (Wolfe & Shank, 2019). Reading Novalis's notes for his unfinished encyclopaedia is a revealing process: in them, he outlines his idea of what is encyclopaedic and what isn't; how knowledge should be presented; what the relations between the sciences are. Bacon explicitly addressed the possibility of future development, and outlines the ethos of the project, so that potential future contributors have something to refer to (Bacon, 1620/1909–14).

The primacy of curation is glaring on Wikipedia, mostly written by anonymous editors, or editors writing under a pseudonym (which is quite close to anonymity). However, its founders are well known and have extensively shared their views about what the project should be. Larry Sanger has also articulated, in his memoir, the reasoning behind a lot of Wikipedia's setup and his opinions on how it developed over time. Jimmy Wales is still the ultimate authority on what Wikipedia should be. At the time of writing, he has just reacted to Elon Musk's floating a purchase of Wikipedia, thereby re-affirming the values of freedom derived from the FOSS movement. It is true that, in practical terms, Sanger and Wales have tried to embody a minimalist approach to curation, consistently with the FOSS ethos: they have left as much space as possible for others to create Wikipedia. Wikipedia was born extremely malleable. But even malleability is a choice with consequences: just like leaving a space without rules means leaving it to those with existing power, those who knew what to do with an unstructured wiki ended up creating the structure of Wikipedia as is today.

Planning an encyclopaedia, and creating the structure that will host future content, is possibly the most meaningful part of the job. Classifying knowledge is an ideological act (Bowker & Star, 1999; Foucault, 1969). Wikipedia shows its roots by organising following the traditional Western classification of knowledge (Van der Velden, 2013). Encyclopaedias can also be designed on the basis of a specific understanding of their use: the 1974 Encyclopaedia Britannica was divided in sections that provided information on the same topics, in increasing levels of depth, thereby operationalising an educational model known as the "circle of learning" (Blair, 2006). I have already described Stallman's plan for a free encyclopaedia, above.

Organising knowledge is not just a matter of translating ideology into a tangible artefact, however: available technologies also influence it. Laying out information based on a circle of learning approach was made possible by the use of cross-references. When encyclopaedias were written on paper, and in alphabetical order, planning was required from the beginning, as volumes would come out in order. Indexing and alphabetical order, in turn, make classification, in some cases, redundant. Wiki technology makes it possible to allow large groups of people to contribute to an encyclopaedia, opening up authorship, and perhaps even redefining the notion of expertise.

Neither technologies nor epistemologies have primacy here: each encyclopaedia is the result of

the interplay of various factors, including their material substrate and the ideology of their creators. An example above all is Conservapedia: despite explicitly presenting itself as the conservative response to Wikipedia, it uses MediaWiki, the wiki software that was developed for and by Wikipedia. This doesn't mean that MediaWiki is neutral, or that it is so infinitely customisable that it can serve any community. It means that, at the very least, it can serve both Conservapedia and Wikipedia.

Wikipedia's genre underdetermines its design: encyclopaedias have been created with objectivity or a specific perspective in mind; using various systems of classification; as courses or reference works. At the same time, though, the encyclopaedic approach to knowledge-production carries certain commitments – the ambition to exhaustiveness, an orientation towards the future – and certain typical means of fulfilling them. Wikipedia has absorbed some characteristics of its predecessors – the specifics are a result of circumstance, material affordances, and ideological influences.

1.5.5. Design choices

In this section, so far, I pursued two, complementary, aims: the first, and most obvious, is to show what ideological, material and social influences have contributed to shape Wikipedia; the second, is to emphasise the range of possibilities that Wikipedia's format affords. Encyclopaedias come in a range of shapes and sizes, but not all shapes and sizes: the encyclopaedic genre, well established by the time of Wikipedia's inception, offered a set of options; Wikipedia's designers selected which to pick, based on their ideological background, social position, personal taste.

Choosing objectivity, in the form of neutrality, as a guiding epistemic principle, is not simply a result of Wikipedia being an encyclopaedia. Rather, it seems to stem from the engagement with Ayn Rand, who considered neutrality as a reachable epistemic ideal, and is the result of the personal positions of Wikipedia's founders. Wikipedia's forkability is rooted in its relationship to the free software movement, where permitting copies is both a safety measure and a result of their ideological commitment to open knowledge. The status of authorship on Wikipedia occupies a space between celebration of the individual, and individual action, expressed in policy (Be bold) and encouraged by the combative environment fostered in community spaces, and (possibly, in some cases, performative) humility – a synthesis of Ayn Rand's celebration of the romantic hero and the traditional unassuming authorship found in encyclopaedias.

An additional player in this field of influences is the materiality of the website itself. Over the course of this dissertation, I will show how design choices have been influenced by the ideological background of Wikipedia's founders, and the collective ethos held by Wikipedians as a community. I will also argue that the reverse is true: the material form Wikipedia takes – a wiki website – presses, in my view, on ideology as well, co-shaping, as well as enforcing, values. Looking closely at the design process shows, I argue, a relational field where choices and objects influence each other – questioning certain traditional understandings of values in design, and the social shaping of technology. In chapter 2, I will outline the theoretical space in which this project exists, and the corresponding methodology used for my investigation.

Conclusion

Wikipedia is a complex machine, where the interactions between people, culture and technologies are varied and idiosyncratic. In order to understand how it works, it is not sufficient to grasp the basic mechanisms by which pages are updated: one needs to appreciate how editing is regulated, what regulation means, and how it intersects with the platform itself.

The minute interdependencies found within Wikipedia challenge widely held understandings of power: the traditional top-down or bottom-up models, where power is held by a specific group of people and wielded over others, makes little sense on Wikipedia. Because of its distributed nature, and of the amplifying role of digital technologies, power hides in unexpected places: for instance, a volunteer programmer who creates a powerful tool for detecting vandalism may have more influence over the project than an administrator who is able to decide whether individuals may or may not participate. Technology can't be framed as the seat of power either: it regulates behaviour in some ways, but it is also created by human coders, whose action is, in turn, afforded and limited by other coded objects. Policies can be evaded by those who know how to, and knowing how to depends, in part, on technological savvy and standing within the community.

Wikipedia needs to be understood, as well, as part of a broader picture: it is an offshoot of hacker culture, a Randian project, and, as an encyclopaedia, it inherits traits from its predecessors. In order to make sense of how Wikipedia works, then, it is important to see how the ethos of its predecessors has seeped into the project, and what the consequences of that inheritance are.

Scholars before me have approached Wikipedia from various angles, looking at its cultural practices, technological affordances, epistemic tenets. Little attention, however, has been dedicated to how the design choices made by those who have developed Wikipedia relate to its epistemology. In the next chapter, I will outline the project on which this dissertation is based, detailing my methodology and the theoretical environment around it.

2. Methodology and theoretical framework

he also had a device which looked rather like a largish electronic calculator. This had about a hundred tiny flat press buttons and a screen about four inches square on which any one of a million "pages" could be summoned at a moment's notice. It looked insanely complicated, and this was one of the reasons why the snug plastic cover it fitted into had the words Don't Panic printed on it in large friendly letters.

Douglas Adams, The Hitchhiker's Guide to the Galaxy

2.1. Introduction

As outlined in chapter 1, this dissertation is a study of how values are embedded in the design of English Wikipedia's website, en.wikipedia.org. In order to conduct my analysis, I have devised a mixed-methods methodology that responds to the specific material and semiotic qualities of digital objects, leveraging tools from applied philosophy of technology – specifically, value-led design approaches and postphenomenology – as well as critical, material approaches to the analysis of digital objects, such as platform studies and critical code studies.

In this chapter, I describe my theoretical framework and methodology, how they relate to each other, and how they have been applied to the generation and analysis of data. I will also trace a history of the research project, which has evolved over time, in part for reasons related to preliminary results, in part to circumstance. Preliminary results revealed how my initial focus on discourse lacked sufficient engagement with the materiality of Wikipedia as a platform; as a consequence, I reframed my approach by understanding Wikipedia's text as a "technotext" (Hayles, 2002), that is: a text that engages with the material substrate on which it is inscribed, which in Wikipedia's case is a digital platform. The circumstantial reason for the change in focus concerns the publication of an important book by a preeminent scholar of Wikipedia, which reached similar conclusions to the ones I envisaged for my own, thereby threatening the novelty of my research.

The necessity to reshape the project pushed my research in unexpected, ultimately more interesting directions, as the methodology described here attests: I have been able to fruitfully combine traditionally unrelated methods, in order to examine a complex digital artefact – its structure, values, and the relations between the two. To the best of my knowledge, this is the first extensive study of this kind, and, I can say with certainty, the first to analyse the entanglements between Wikipedia's ideology, culture, and design.

2.2. Research questions

It is important to note that the focus and scope of my doctoral project has changed over time, taking a sharp turn around November 2022. Below, I will outline the earlier iteration of the project,

the reasons for the change of direction, and the final research questions that this thesis addresses directly.

As the project, and consequently its methodology, has evolved over time, it is necessary to provide a brief overview of the various stages of development. At the inception of my doctoral studies, I was interested in Wikipedia as a case study for the investigation of the dialogical construction of knowledge. Inspired by theories of collective knowledge-building geared around dialogues, such as the accounts provided by Hugo Mercier (Mercier, 2016; Mercier & Sperber, 2011), Catalina Dutilh-Novaes (Novaes, 2015, 2018, 2020b, 2020a) and Sinan Dogramaci (Dogramaci, 2015). I was particularly drawn to Wikipedia for multiple reasons: first, it is a large, and successful, collective knowledge-production endeavour; secondly, it explicitly used deliberation as a tool for knowledge-production, in the form of discussions on talk pages; third, access to deliberation is easy and complete: all deliberation occurs on talk pages, which are publicly available, providing easily accessible, relevant and abundant data on how editors make decisions on content. The initial, exploratory coding scheme reflects these priorities; I will sketch its contours in section 2.6, but won't go into detail, as it is mostly irrelevant to the current state of the project.

The project changed as I learnt more about Wikipedia. As boyd (2008) argues, since online communities are not bounded by geography, but rather defined by a network that doesn't correspond to an existing culture, they can't be understood within a pre-existing context of norms. Rather, the researcher ought to understand them in their own terms, calling for an iterative methodology that alternates gathering information on norms and modifying the methodology as needed, to start the process again (boyd, 2008). The boundaries of any research project are socially constructed – interaction with the community being researched is part and parcel of defining the scope of research (*ibid.*).

I came to Wikipedia as a novice, much like the philosophers whose work I described in 1.4.2, and as such I was not aware of the specific dynamics that inform knowledge-production on Wikipedia. Over time, I realised how Wikipedia cannot simply be understood as a site of dialogical construction of knowledge, ignoring the social context within deliberation takes place: to do so wouldn't simply constitute an oversimplification – rather, as I argued in section 1.4.2, it would lead to a misleading representation of affairs. Consequently, my project took a more sociological flavour, turning into an investigation in social epistemology. At this stage, my research questions were:

1. How is knowledge produced on Wikipedia?
 - a. How do Wikipedia's interface, policy, and community practice influence the process of knowledge-production?
 - b. How is expertise constructed on Wikipedia?
 - c. How are decisions about Wikipedia's content made?
2. What values are designed into Wikipedia's interface, encoded into its policy, and shared within its community?

Following Ford, I framed Wikipedia as a sociotechnical system, and set out to investigate how it produces knowledge, and what kind of epistemology was embedded in said process. I was particularly interested in what expertise means within a system that is built to ignore expertise as traditionally construed in adjacent epistemic environments (e.g. academia), and what the mechanisms that defined content production – the closest one can get to a Wikipedia conception of truth. I planned to leverage concepts from philosophy of technology in order to make sense of the interaction between Wikipedians and the platform, to analyse talk pages in terms of what arguments Wikipedians appeal to in deliberation, under the assumption that those principles that went unchallenged constituted the collectively understood core of the community’s epistemic values, and to interview Wikipedians and Wikimedia Foundation staff in order to explicitly ask what those core values were, and how exactly they were conceptualised in practice. Policy analysis provided context but, as explained in chapter 1, it can’t be taken at face value: engagement with the community would allow me to understand the relationship between practice and policy, as well as what parts of policy are ignored versus genuinely enacted in practice. At the same time, I was gathering technical documentation. I performed a thematic analysis of the data, which is available in Appendix A. In parallel, I was drawing connections between my partial findings and the cultural and intellectual influences on Wikipedia’s design: Ayn Rand’s philosophy and the ethos of the FOSS movement. The preliminary analysis conducted in this context has informed the further development of the project.

Broadly speaking, my work was meant to build on, and complement, work by other authors in the same space, primarily Heather Ford and Stuart R. Geiger’s. In November 2022, the book *Writing the revolution* was published, in which Heather Ford describes the process of knowledge-production on Wikipedia, by focusing on how a single event, the Egyptian revolution of the early 2010s, also known as the Arab Spring, was covered by Wikipedians. In line with her previous work, she used a sociotechnical perspective, describing how facts are constructed through interaction with the platform. In her conclusion, she draws lessons around epistemic justice, engaging with the materiality of Wikipedia’s platform, and how it influences, through affordances, the knowledge-production process. In other words, Ford addressed, robustly and mostly satisfactorily, my Research question 1, including all of the sub-questions. The gap I had identified that I was trying to address with my research, was rapidly closing.

While I had sufficient data to elaborate my own account, it became clear that my own conclusions would rest heavily on Ford’s, reducing the likelihood of a novel contribution. However, Ford left question 2, about the construction of values on Wikipedia, essentially untouched – and so did the rest of the literature. Concurrently, I noticed an interesting continuity across social and technical data, between values that were culturally relevant to the Wikipedian community, and ways in which the platform was designed. Given the central role of the concept of neutrality, which was confirmed by my interviewees and frequently appealed to in talk pages, I had ended up probing how it manifested in different domains: epistemic, aesthetic, affective, ethical. It seemed interesting to pivot the focus slightly, and work more precisely on design, with a view to work within, and contribute to, the field of values in design, which led the current iteration of the project, outlined in the following. I selected the body of data that was more specifically geared around design, and sought to add to it, guided by a new set of questions, which, in essence, articulate part of initial research

question 2:

1. How is Wikipedia's epistemology embedded in the design of Wikipedia's website?
2. How have the major ideological influences on Wikipedia's culture impacted the design of the knowledge-production process on Wikipedia?
3. What theoretical contribution to the field of philosophy of technology can be gained from a close reading of en.wikipedia.org?

The analysis of data around design and aesthetics took the form of a second thematic analysis, detailed in Appendix B. A complete breakdown of the body of data used in this project can be found in section 2.6. As an unexpected result, looking at Wikipedia's infrastructure also allowed me to reach important conclusions about Wikipedia's social epistemology, as I will argue in chapters 3 and 4. As the project evolved, some theoretical pillars remained, while others changed with the nature of the questions to be answered and data to be analysed. In the next section, I will describe the theoretical framework that informs the argument presented in this thesis. I will then move on to methodology and the technical details of my research.

2.3. Theoretical framework

In this section I describe my theoretical framework, which is a combination of philosophy of technology and design, in the form of postphenomenology and value-led design approaches, and critical approaches to coded artefacts, specifically critical code studies and platforms studies. The former concern themselves with the role fulfilled by technologies within human existence: postphenomenology focuses on the way specific technologies place themselves in relation to humans, and mediate relationships between humans and their surroundings, as well as between humans themselves; value-led design approaches take a more proactive attitude, in building a theoretically sound methodology aimed at deliberately creating technology that is, at the very least, ethically sound, if not, more ambitiously, promoting human flourishing by embedding desirable values. The latter – critical code studies and platform studies – study cultural influences on digital objects, at different levels.

My aim, in bringing these two methodologies together, was to examine how ideology and technics influence each other. I had observed how Wikipedia's design was discussed in ethically loaded terms, and how technical or epistemic values, such as, for instance, efficiency or neutrality, were deployed rhetorically as human values, to support arguments. I was interested in tracing the path values take across the social and technical realms. I integrate the framework above by drawing on the work of Maurice Merleau-Ponty, Michel Foucault, and Gilbert Simondon, bearing on, respectively, the embodiment of subjectivity, the shaping of the self, and the relationality in the design process.

2.3.1. Postphenomenology

The term “postphenomenology” labels a cluster of several theories, which interact with each other and form a broadly coherent approach to the philosophy of technology, mainly aimed at creating accounts of specific devices. Its roots can be traced back to classic phenomenology, a philosophical approach centred on the investigation of human experience. Distinctively,

postphenomenology focuses on relationships between machines, humans and the world, and is committed to grounding analysis in empirical research, examining concrete, specific artefacts (Verbeek, 2005). According to postphenomenologists, machines play an active role in constituting the relations they are a part of (Rosenberger & Verbeek, 2015). Machines invite users to act in specific ways, and they do so in virtue of their design. Since the relation between the user and the machine is formed in interaction, the machine itself contributes to it. For instance, in presence of an anthropomorphic robot, I will lean towards treating it like a quasi-person: I might say hello to it. The robot itself, by looking similar to a human, taps into my inclination to interpret faces as fellow human beings, thereby contributing to making our relation an “alterity relation” (Verbeek, 2005).

Studies in postphenomenology vary in terms of the methodology adopted; I sum up below the main points of postphenomenological methods, according to Rosenberger and Verbeek (2015).

1. Postphenomenological studies have an empirical basis: whether cited from the literature or based on directly acquired data, research in this field emerges from observation. This aligns with the focus on, and relevance given to, specific technologies: postphenomenology doesn't apply a prefabricated interpretation to technology; it learns from technology in order to enrich the theory. Of course, observations are made within a theoretical framework, and some tools are readily available to be used. However, these are conceived as a starting point, useful to make sense of observations before proceeding further.
2. A basic step in examining technologies is identifying how they shape relations between themselves, humans and the world.
3. Once the relations have been identified, an account is developed of how “a specific ‘world’ is constituted, as well as a specific ‘subject’” (p. 31).
4. Based on the elements above, an analysis is made of how the technology in question contributes to shaping one or more areas of human experience (the production of knowledge, the appreciation of art, decision-making in a biomedical setting, etc.), along one or more dimensions (ethical, epistemological, political, etc.). It is worth pointing out here that postphenomenology doesn't hold a deterministic stance: technology doesn't entirely shape behaviour or alienate people from themselves. It can “invite” (Ihde, 1990), “evoke, provoke” (Verbeek, 2005) a way of being used, which is context-sensitive and escapable, while still exerting a powerful influence. In most circumstances, users approach machines as devices that will improve their experience or simplify an operation and go along with the use suggested by the object in a given context. In other words, convenience nudges the subject's behaviour towards compliance with the object's “implicit user's manual” (Procee, cit. in Verbeek, 2005, p. 115).

I have chosen postphenomenology for this project because it offers useful tools to think about

specific artifacts, rather than technology in general, and a conception of technology that allows for changing the way we design and relate to it (Feenberg, 1999). Postphenomenology also allows conceptual space for the analysis of embodiment.

In order to understand technologies in a way that allows for change, a central concept in postphenomenology is that of “black box”, developed by Bruno Latour in the context of the sociology of science, and absorbed in postphenomenology through the work of Peter-Paul Verbeek and especially Andrew Feenberg. When a technology is being developed and introduced to a new context, the actors who influence it (designers, politicians, customers, and so forth) are apparent to all; at some point, through an operation called closure, this network of social influences is forgotten, and only the object remains (Feenberg, 1999). Opening the black box brings to light its hidden meanings. This can happen, for instance, when the object breaks, and it has to be physically taken apart for fixing, revealing the materials that constitute it, each with its own history (Feenberg, 1999); but also its surviving dependence on actors involved in its creation, as for instance in the case of a television that has to be sent back to the manufacturer for repair (Verbeek, 2005); and the dependence of humans on its functioning, for instance through the disruption – and frustration – created by a malfunctioning computer (Verbeek, 2005). The knowledge acquired by opening the black box can allow us to put forward suggestions for reform, including a democratisation of some parts of the process (Feenberg, 1999).

Embodiment is addressed by the concept of multistability, articulated by Don Ihde in *Technology and the Lifeworld* (1990). According to Ihde, an artefact can assume different meanings depending on users (Ihde, 1990). His focus is on culture: for instance, the same object can be seen as a tool by some, and a decorative item by others. In *Technology and the Lifeworld*, he describes how in the 1930s, oval sardine tins left in New Guinea by Australians were incorporated in headwear as decoration, where similarly shaped shells were customarily positioned (Ihde, 1990). Verbeek has abstracted the notion of multistability from the cultural context, picking up on its kernel – the fact that technologies are constituted in interaction – and generalising it. Multistability, in Verbeek, becomes both broader and narrower: broader, because it applies beyond cultural differences; narrower, because it ends up ignoring the context in which the subject acts.

Within the context of Wikipedia, the role of multistability is heightened by the extensive power that the user community has on the technology itself. Not only, as in traditional multistability, can users co-construct the technology through use – on Wikipedia, users can even participate in the programming of the platform. Co-constitution through use takes the form of, for instance, *edit warring* (Wikipedia: Edit warring), which means reverting someone’s edits over and over again in order to make a point. This way editing facilities are turned, through (mis)use, into sites of conflict: the basic ability of editing the text and reverting edits made by others, which is the core of editors’ activity, has been weaponised. The participation of Wikipedians in shaping the platform reaches even deeper: Wikipedians intervene, for instance, through the creation of bots; they cooperate with the professional programmers working for the Wikimedia Foundation, which operates Wikipedia; and can insert of new features in the very fabric of the wiki, such as userboxes.

2.3.2. Embodiment and the shaping of the self: Merleau-Ponty and Foucault

My application of postphenomenology is based on a deeper engagement with the specific embodiment of the subject. Rather than dealing with the interaction between Wikipedia and a generic subject, as in Verbeek, or between Wikipedia and a culturally situated, but not fully articulated, subject, I wish to open up the possibility of investigating the subject as embodied.

In order to do so, I leverage Merleau-Ponty's account of subjectivity and Michel Foucault's concept of "technologies of the self", which I will now explain in turn. Merleau-Ponty's account is distinctive in its identification of the subject with a body. The term "body", in his work, doesn't just refer to the physical body, but also to those material, yet not simply anatomical, features which characterise embodied beings and mark them as individuals: their location in space, cultural context, ancestors. The body is situated within, an inextricable from, a geographical, cultural and historical context. The body provides the coordinates to place an individual in time and space, within their community and their own experience. I use "body" to signal this material engagement with the environment, as opposed to identifying individuals with pure minds, or in isolation from their contexts.

The way the body-subject exists in the world rests on the habitual body, and on context. The habitual body is constituted by the sedimentation of experience:

The habitual body can act as guarantee for the body [...]How can I perceive objects as manipulatable when I can no longer manipulate them? The manipulatable must have ceased to be what I am now manipulating, and become what one can manipulate; it must have ceased to be a thing manipulatable for me and become a thing manipulatable in itself. Correspondingly, my body must be apprehended not only in an experience which is instantaneous, peculiar to itself and complete in itself, but also in some general aspect and in the light of an impersonal being.

(Merleau-Ponty, 1945/2005)

One of his examples is that of a woman habitually wearing a large hat: in order not to bump into furniture and walls, she has to move in a particular way that accommodates the presence of the hat on her head. What happens then, is that she is so used to moving in a way that accommodates the hat, that she moves in the same way even when she is not wearing it, essentially making up for an imaginary hat. For instance, she might walk through a door sideways. The habitual body, then, influences what became later known as "affordances", the invitations that we receive from object to use them in a given way. A door affords walking through. A wall doesn't. In Merleau-Ponty's words, the habitual body is the foundation for the "I-can" – the understanding of the world as composed of things that can and can't be done by me. What is manipulatable (the "I-can" in Merleau-Ponty's terminology) depends on one's habitual body. This space between an object that invites and a user who is invited is where Verbeek's theory applies. Extrapolating from this notion, a digital environment can be thought of in terms of embodiment, as it is constituted by the habitual body: a digital environment is manipulatable based on one's own experience of similar environments. Since

digital environments, like all artefacts, are multistable, the role they assume in interaction varies slightly depending on who is experiencing them.

In Merleau-Ponty, the body is also historical: it is influenced by the historical moment, context, geography, ways in which one is seen by others (Merleau-Ponty, 1945/2005). In his view, an “intentional arc” backs up the “life of consciousness” (which, in Merleau-Ponty, is indistinguishable from the body), by projecting our personal history (“our past, our future”), our positionality (“our human setting, our physical (...) situation”) as well as our beliefs (“our ideological and moral situation”). In sum, “our being situated in all these respects” (all quotes in this paragraph, p. 157): the situatedness of the body informs the behaviour of the subject. It is important to note that such situatedness is not deterministic: we don’t act the way we act exclusively because of our position within our personal coordinates. In Merleau-Ponty’s view, history, psychology and physiology are all part of a sole movement, directed towards the world. They resist the reduction to one another: our position is “not a destiny, (...) just as clothing, jewellery and love transfigure the biological needs from which they arise” (p. 101).

Considering situatedness as an important aspect of the encounter between an artefact and a user is an enrichment of Don Ihde’s concept of multistability. By applying Merleau-Ponty’s view, I am expanding the basis on which the notion of multistability rests: it is not just culture that turns artefacts into different technologies depending on whose hands they are in. It is also gender, race, ability, class, and possibly more sets of coordinates that might even vary depending on technology.

Another way in which the self is shaped through action is via mechanisms termed by Foucault “Technologies of the self”. In Foucault’s work, technologies of the self are practices, enforced through interactions with others, aimed at shaping the subject (Foucault, 1994). The meaning Foucault attributes to the word “technology” in “technologies of the self” doesn’t necessarily align with common understandings of technologies, and indeed I am not claiming that Wikipedia’s interface is a technology of the self. What I am claiming follows from the way Foucault theorised the functioning of the technologies of the self. After defining four types of technologies, some of which, e.g. technologies of production, will sit more comfortably within widely accepted definition of technology, he adds:

Each [type of technology] implies certain modes of training and modification of individuals, not only in the obvious sense of acquiring certain skills but also in the sense of acquiring certain attitudes. I wanted to show both their specific nature and their constant interaction. For instance, the relation between manipulating things and domination appears clearly in Karl Marx’s *Capital*, where every technique of production requires modification of individual conduct – not only skills but also attitudes.

(Foucault, 1994, p. 225)

Hence, according to Foucault, participating to a process of production trains the self in ways that reach beyond learning: individuals also acquire behaviours, and a certain way of being in the world. The relevance of this concept to Wikipedia in particular, and sociotechnical systems in general, is that individuals don’t simply participate in processes of production: they are a part of the production process, and inevitably shaped by it, at the same time as they shape the process itself. I will expand

on this topic in chapters 4 and 5.

2.3.3. Value-led design theories

Since the 1990s, a series of theories have emerged that concern themselves with the way values can be designed into technology. Scholarship within the approaches known as value sensitive design (Friedman, Kahn, & Borning, 2002), design for values (Van de Poel, 2015), values in design (Bednar & Spiekermann, 2022), and value-oriented engineering (Longo *et al.*, 2020) describe methodologies by which specific moral values can be embedded in artefacts, as well as theorising the way in which values exist, interact, and change through artefacts. These theories overlap widely and are very close to each other, in ways that don't have a bearing on my work, so I will frequently cluster them together under the label "value-led design".

Value-led design accounts of embedding are practice-oriented: they provide a structured methodology for the deliberate embedding of values in technological artefacts. There are local differences, but broadly speaking, value-led design approaches propose a three-phase model of embedding. First, designers define the values to be embedded into the product: values are described as being elicited (Bednar & Spiekermann, 2022; Friedman & Hendry, 2019; Van de Poel, 2015; Witteman *et al.*, 2016), or discovered (Flanagan *et al.*, 2008). The second step consists in settling value tensions or trade-offs, in order to come up with a final list of values (Friedman & Hendry, 2019; Miller *et al.*, 2007; Van de Kaa *et al.*, 2020; Van de Poel, 2021). Finally, values can be operationalised, or transformed into design specifications (Flanagan *et al.*, 2008; Friedman & Hendry, 2019).

Since value sensitive design is the most systematically theorised approach, I have adopted their theoretical principles in guiding my analysis. Based on Friedman & Hendry (2019), the principles most salient to my research can be summarised as follows:

1. The object of analysis of value sensitive design is defined as tools, defined as extensions of human action; technology, defined as tools that leverage scientific knowledge²; and infrastructure, defined as the material and organisational means facilitating the operations of a collective body. In this sense, then, the object of analysis is not simply artefacts: a management system falls under the remit of value sensitive design as well.

² This definition of technology, much like all other definitions of technology, is not obvious, nor uncontroversial. Some might argue that designing technology doesn't require scientific knowledge (and I would agree); others that scientific knowledge is constructed through technology (I would also agree), or even that scientific knowledge is a technology (I would suspend judgment on this point). In fairness to Friedman & Hendry, they do acknowledge a continuity between technology, tools and infrastructure, partly on the grounds of the contested role of knowledge in designing technology. Regardless, settling the dispute over the relation between science and technology is, for present purposes, and in methodological terms, irrelevant. The key point here, that informs my methodology, is that an analysis in value sensitive design terms can be conducted of a variety of entities, all of which have two things in common: they provide a structure for human action, and come to exist in their specific form for a specific purpose.

2. Values are “what is important to people in their lives, with a focus on ethics and morality” (p.24)
3. Values don’t exist in isolation, but interact with each other.
4. Interactional stance: Values can be embodied in technology, through design – however, people use and modify artefacts according to their values, so understanding values in technology means understanding the context of use and production more widely.

In the context of this project, I have applied and, in part, adapted the principles above:

1. In the case of Wikipedia, my analysis is not limited to technical artefacts, but extends to the process of knowledge-production that is implemented through Wikipedia’s digital platform.
2. I take a generous approach to values, and include aesthetic and instrumental values as well, thereby expanding the territory to kinds of values investigated by other approaches.
3. Within the value-led design literature, interactions between values are understood in terms of trade-offs, that occur either at the point of design, where there might be conflict over what values should be implemented, or at the point of use, where users may modify or use artefacts based on their own values. I take a broader view of interaction, and see values as interacting through material means as well – that is, through trade-offs in operationalisation.

It is worth noting there that one of the outcomes of the research presented in this dissertation is a theoretical contribution to the area of values in design, which leverages empirical insight on Wikipedia to enrich existing accounts (see chapter 5). In order to enrich our understanding of the role of values within design processes, it would be useful, in my view, to adopt a relational attitude. I will develop this argument in chapter 5; at this juncture, however, it is helpful to lay the foundations by engaging with a key issue: hylomorphism.

2.3.4. Relationality during the design process: Simondon

As I will argue in chapter 5, relational dimensions of the design process are underexplored in value-led design theories, and theories of the value-ladenness of technology more broadly. Significant steps towards understanding relationality in design were made by philosopher Gilbert Simondon, whose critique of hylomorphism informs my understanding of value-ladenness. Simondon addresses hylomorphism in *L’Individu et sa genèse* (1966). The term “hylomorphism”, generally speaking, refers to the Aristotelian doctrine stipulating that objects are made of matter and form (Ainsworth, 2024). Simondon applies it to the making of objects, defining hylomorphism, in this context, as the

idea that individuation – the coming to existence of a specific, contingent entity – occurs as a result of some kind of undifferentiated matter being shaped through forms. So the two elements here are: the forms – some kind of mould that is imposed on matter – and raw matter.

Simondon moves two main objections to hylomorphism: one historical, one conceptual. The historical point is that, according to his account, the reason of our misrepresentation of the design and creation process is social: by attributing humans power over matter, we reproduce the attitude that masters have over slaves, framing themselves as those who provide forms, through orders, and slaves as those who operate at a material level, translating orders into objects (Chabot, 2013).

Additionally, Simondon argues, hylomorphism is conceptually shaky in that it assumes that matter is undifferentiated – a necessary precondition for the human ability to mould it into any desired form. According to Simondon, on the contrary, both forms and raw matter are, in and of themselves, constructed and already somewhat differentiated. In order to make his point, he describes the making of a brick. One could say, hylomorphically, that the brick is made out of a raw material, clay, shaped by pressing it into a mould (the form). However, Simondon writes, that's not an entirely accurate description:

Shaping clay doesn't mean imposing a cuboid shape to raw clay: it means pressing prepared clay into a manufactured mould. If we start from the two ends of the technological chain – the cuboid, and the clay in the quarry – one has the impression of fostering, in the technical operation, an encounter between two realities of heterogeneous domains, and of establishing a mediation, through communication, between an interelement, macrophysical order, bigger than the individual, and an intraelement, microphysical order, smaller than the individual.

(L'Individu et sa genèse, p.38, my translation)

At the level of the macrophysical order then, Simondon observes that the mould is a technology in and of itself: it is made to shape bricks. The matter itself, the clay, is also prepared, and it is specifically prepared to easily adapt to the mould: it is coarse enough that it will dry in shape – thereby tending to the final purpose of the operation: making a brick – but loose enough to slip into the corners of the mould. Therefore, the operation of moulding a brick retroactively shapes the construction of both form and matter.

Looking at the microphysical order, Simondon notes how materials present internal features, that allow or inhibit certain kinds of usage: he calls these internal features “implicit forms”. Wood, for example, writes Simondon, contains an internal structure that is variable and can lend itself to different uses, depending on species, or even on individual pieces of wood. Further, each tree is not internally uniform: it contains fibres and knots. One can decide to cut the wood and disregard these internal structures, or, as Simondon puts it, impose explicit forms on the wood, but that comes at a cost: going against the grain is likely to result in a more fragile final product. The better way of cutting wood consists in using tools that allow to follow and leverage its internal structures. Implicit forms appear to work through the mechanism that we now call affordances : they constitute the mechanism by which a material lends itself to certain uses, and resists others.

It should be noted that Simondon's position on technology can be read as determinist, even though there is some debate on this point (see Feenberg, 2017, who tries to reconcile Simondon's

account with a constructivist perspective). My personal view is that Simondon was incoherent on this point: perhaps because of the time of writing, when a full debate over the position of technology in society hadn't developed yet, he might have thought he didn't need to position himself towards this specific question, and/or didn't think of tending to the inconsistencies within his own body of work. Regardless, his work on individuation, which I am mostly referencing here, is in my view compatible with a relational, non-deterministic account of technology.

2.3.5. Platform studies

My work is also informed by the fields of critical code studies and platform studies. Both fields borrow from semiotics and aesthetics in order to analyse coded objects, understood in the context of their creation and use. The following quote applies to critical code studies, but can be said of platform studies as well:

Critical code studies is an approach to code studies that applies critical hermeneutics to the interpretation of computer code, program architecture, and documentation within a sociohistorical context. CCS holds that the lines of code of a program are not value-neutral and can be analyzed using the theoretical approaches applied to other semiotic systems, in addition to particular interpretive methods developed specifically for the discussions of programs.

(Marino, 2020, p. 49)

The guiding principle of critical code studies and platforms studies is that code, and Web platforms, are cultural objects, similarly to works of art and everyday objects. They can be read, understood, and evaluated from an ethical point of view in the context of their creation, in terms of their aesthetic properties, and in terms of what they do.

Code is written to instruct a machine, but computers don't require symbolic representation – humans do (Marino, 2020; Montfort & Bogost, 2009). Hence, any semiotically salient trait of code – for instance, the fact that it's written in words, modelled on the human English language, as opposed to equally computer-friendly gibberish – has to do with its audiences (other coders, the public, management) and the culture of those who create it (Montfort & Bogost, 2009). The cultural situatedness of code becomes apparent when looking at cases in which code is written with specific semiotic characteristics: for instance, obfuscated code and weird languages, both aimed at making it difficult for others to read code; if code were written only as a set of instructions for a machine, it would be meaningless to speak of obfuscation (Mateas & Montfort, 2005).

Code is as much a set of instructions as it is a means of communication between humans. Even in non-limit situations such as obfuscated code, programs carry traces of their authors, in ways that can have important consequences. The same result, for instance, can be obtained by coding with efficiency in mind, hence using fewer lines of code, or more diffusely. At first brush, one might think the two are equivalent, or, if we are able to do the same thing with less use of resources, it might be better to code more efficiently. Now, coding efficiently also means coding opaquely: it would be harder, for coders who are not the author, to understand what the thinking was behind that specific chunk of code, and consequently it might be harder to change it, improve on it, or find issues (bugs) if needed (Marino, 2020). By contrast, extended code, if well-documented, lends itself to be

understood and modified by others: in other words, it is more collaborative (Marino, 2020). It is easy to see the ethical colouring of this difference: efficient code is individualist; more diffuse, well-documented code, is collectivist.

This becomes even clearer when thinking about the cultural context around coding. First, different ways of coding are gendered: in the example above, women are more likely to write code diffusely but collaboratively than men, who tend to code efficiently as a sign of prowess (Marino, 2020), since efficiency is prized in competitive, masculinist coding contexts (Ensmenger, 2015). Originally, writing short code fulfilled the practical purpose of using scarce computational resources wisely; writing compact code, however, soon became a way to demonstrate skill, performed in competitive contexts (Ensmenger, 2015). Secondly, open source software – that is, as explained in chapter 1, programs that are supposed to be modifiable and reusable under specific circumstances – should be written to allow easy modifications; therefore, not cryptically. The efficiency example also serves to illustrate what is meant by aesthetics of code: efficient code is considered beautiful in coding communities (Coleman, 2013; Levy, 1994), perhaps due to the connection with technical ability explained above.

Considering the materiality of digital objects, finally, means considering what they can do, and the relationship between their physical existence and the content they display. Platforms, understood as structures that make it possible to create digital objects, influence the way said objects are created, by supporting certain functions, and not others (Montfort & Bogost, 2009). In some cases, the influence can be direct, for instance when technical limitations are concerned; in others, the influence is more subtle, and concerns the interplay of various elements of the platform itself (Montfort & Bogost, 2009). Coded infrastructure empowers and constrains human action as well: Internet architecture can be understood as material regulation (Galloway, 2001; Lessig, 2006).

In Wikipedia's case, a further semiotic consideration is necessary. Wikipedia, as well as being a platform, is built to produce, and displays, text (as well as images, and other kinds of media – here I will limit myself to text). Taking materiality seriously also impacts the analysis of content, which has to be understood in relation to the substrate that hosts it, as conceptualised in the notion of “technotexts” (Hayles, 2002). In order to understand and critique text, it is not sufficient to understand its linguistic meaning: text needs to be understood in the context of the platform that hosts it. In other words, in order to understand text, we need to be aware of its substrate: is it printed on paper, displayed on a screen, or etched in clay? How does the substrate change our experience of reading it? Do the writers keep the substrate into consideration when writing? When looking at hypertext, for instance, one needs to consider the multiplicity of paths allowed by the form: the text, by virtue of links, is not linear; it sprawls in different directions, and allows for a different relationship between the reader and the author. Hence, when writing hypertext, one might feel freed from the duty of exhaustiveness; rather, be inclined to be brief, make a specific point, and link to in-depth, but not immediately salient, information stored elsewhere.

Reading Web copy as technotexts – asking who wrote it and why, how it relates to the platform, and so forth – allows to see deeper into the complex relationships between text and materiality. Moschini (2018), for instance, developed a five-layered approach to the analysis of social media platforms, which encompasses an analysis of the content (text), the interface, the flow of information, the organisation of information, and the architecture of the platform. Such a system allows to

integrate the semiotic and material properties of social media.

In Wikipedia's case, different texts perform different functions, intertwined with the materiality of the platform: for instance, policies are different from articles, in that they calcify a practice, thereby acquiring a semi-prescriptive function. Part of the authoritativeness of policies is that they can't be easily modified; their rigidity is both embedded in the platform, by making the page itself unmodifiable by most, and expressed through it, by displaying a lock next to the title of the page. In turn, the way Wikipedia's coded infrastructure is designed is regulated by policies, and influenced by social and cultural norms, some of which are constructed discursively in spaces regulated by code, such as forums. The entanglements between culture, software, and text, then, are complex and can be read simultaneously in multiple directions. My analysis attempts to reckon with these intricate dynamics.

I am not the first person to take this approach to Wikipedia. Heather Ford, in her work about Wikipedia, has engaged in forms of platform analysis, albeit without characterising them as such (Ford, 2022), and, crucially, limiting her analysis to how certain coded objects work, without considering the context of their production. I aim to delve deeper into the cultural meaning of Wikipedia's platform.

2.4. Data collection

My analysis is based on observations conducted over the course of five years, and the triangulation of various kinds of data: written documents, interviews, and conversations between members of the Wikipedia community. My intention was to trace how values take different forms based on where they are constructed and expressed, but – I assumed – maintaining a recognisable core form, that would allow to identify them. One can think of the various spaces that can embed and construct values as substrates: written documents, such as policies and strategy papers, are one such substrate; so are human agents (contributors, Wikimedia Foundation employees, coders) and digital artefacts. In this section, I describe how I approached documents, participants and infrastructure.

The description of these various research activities reflects the evolution of the project: some questions I asked Wikipedians to begin with, for instance, are seemingly irrelevant, in the sense that I wouldn't ask them if the initial design were to study how values are embedded in Wikipedia. I include them for completeness, and because they provided helpful context in understanding Wikipedia's culture.

2.4.1. Approaching Wikipedia

Before starting this project, I had never engaged with Wikipedia as an editor. I created an account on 20 August 2019, in order to register my presence as a researcher (this is considered good practice within the Wikipedia community, see section 1.6.1 for details), and communicate with editors. I then received an automated message, which directed me to "WP:WELCOME", and followed links to the policies and guidelines that were indicated as fundamental. I then proceeded to read community pages, guidelines and other documentation as the project unfolded.

In order to acquire an embodied understanding of Wikipedia's process, I have also edited two pages, following my own interests and availability of information. I decided to leverage my native

knowledge of Italian to add information not available in English to the “Bolognese sauce” article (Fig. 2.1), and my philosophical background to add a description of postphenomenology to the page dedicated to philosopher Don Ihde (Fig. 2.2).

Of course, editing a couple of pages constituted an extremely limited understanding of the experience of editing Wikipedia. In order to understand the community, I used existing ethnographies of Wikipedia (Jemelniak, 2014; Lih, 2009) and a guide to editing (Ayers, Matthews, & Yates, 2008), all written by Wikipedian authors. Another useful source of information on Wikipedia’s culture is constituted by essays, which are documents written by Wikipedians

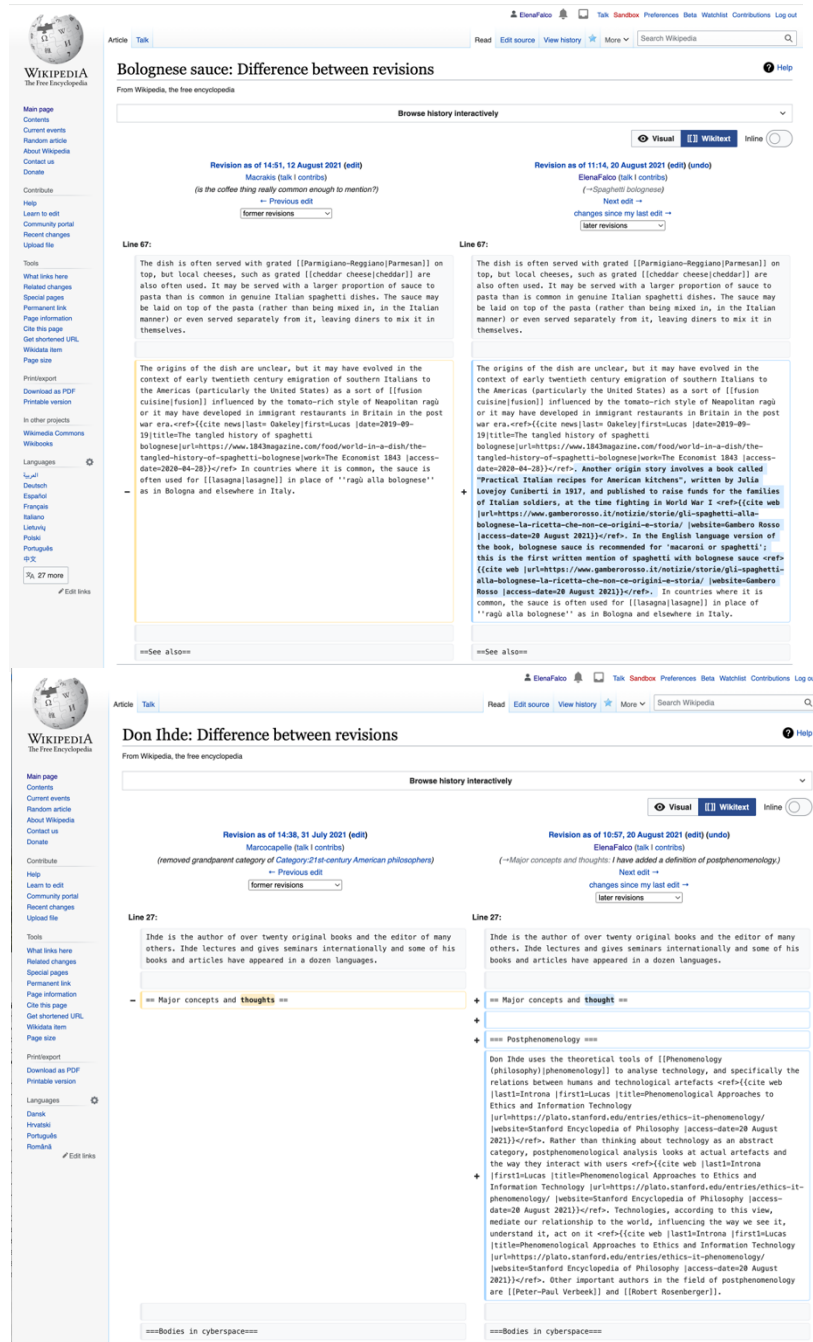


Figure 2.2.1. My edit on the page "Don Ihde". Screenshot taken 1 February 2024.

themselves and published on Wikipedia. Essays cover a wide range of topics, from the application of policies to reflections on the functioning of Wikipedia. This step has been useful as a framework to orient myself within the Wikipedia space: being comfortable with the basic rules and coordinating was necessary to understand what was going on. Knowing policy also allowed me to make sense of how Wikipedians themselves were behaving in talk pages: whether they were following the rules or not, and if they were, whether literally or by embracing their spirit.

Finally, writings by Wikipedia co-founder Larry Sanger are a very useful source of information. Larry Sanger has an academic background in epistemology, and was the creator of some fundamental policies. He has reflected publicly on his role in Wikipedia, by publishing a memoir of his work on the encyclopaedia, several essays, and giving interviews about his Wikipedia. This corpus contains useful insight on the motivations of the founders, as well as explanations of the spirit in which policy was written. Larry Sanger hasn't been a member of the Wikipedia community for a long time – his first-hand knowledge pertains the inception of the project. While of course policy has evolved since the inception of the project, Wikipedia is a fairly conservative project, hence it is reasonable to use Sanger's reflections in order to interpret current community practice.

2.4.2. Conversations with Wikipedians

Over the course of my investigation, I have had various conversations with people connected to Wikipedia, including volunteer editors and staff. Some conversations have been informal, taking place on our user talk pages – the spaces where Wikipedians communicate publicly, by editing each other's profile – while others have taken the form of interviews and email exchanges. In total, I have contacted 33 people and a mailing list. I have then conducted one exploratory interview with a Wikimedia Foundation staff member, and interviewed five more people – three Wikimedia Foundation employees, and two volunteer Wikipedians. The conversations with volunteers have been extensive: I have exchanged long emails with both, and had a continuous engagement with one volunteer, who is interested in epistemology and has also read an early draft of a paper I published on the subject of Wikipedia. Participants have been identified partly through participation to conversations I was interested to, on forums or talk pages, partly via snowballing.

All interviews were semi-structured. Having been designed to answer the first set of research questions, interview questions focus on the experience of editing Wikipedia, and on debate as a tool for knowledge-production. However, from the beginning I factored in a high level of flexibility, that would allow me to adapt to the role of the interviewee. Whenever a participant seemed keen to talk about a specific aspect of their practice, or to express views about a specific aspect of Wikipedia, I have allowed them to drive the conversation; often, participants would spontaneously talk about values, which alerted me to their importance, and in retrospect, provided me with valuable insight into Wikipedia's value system. The more free-flowing part of interviews has, then, borne fruit in terms of the final set of questions as well.

For the first round, I started with a Wikimedia employee and community organiser, and, proceeding by snowballing, I moved forward to interview others. Initial questions for this round were:

- Do you edit Wikipedia? [If yes] Did you start before or after you joined the Foundation?
- What's your day-to-day job?
- Do you any advice on how to interact respectfully with the volunteer community?

- How much interaction is there between the Foundation and the volunteer community?
- What is your personal relationship with the volunteer community?

For the second round, I interviewed volunteer members of the community, to address questions about the knowledge-production process, how they felt about the environment on talk pages (thereby providing information on the role of affect on Wikipedia) and what kind of relationship they have with other Wikipedians. Initial set of questions:

- Do you still edit Wikipedia?
- Talk me through a standard editing session. How often do you edit? When do you edit (at work, in the evenings, in breaks, etc.)? What's your setup (are you on your phone, at your desk on your laptop, etc.) How do you decide what to do, specifically?
- How did you find editing Wikipedia when you first started? How do you find it now?
- Why do you edit Wikipedia?
- What is your relationship with other Wikipedians?
- How often do you interact with them?
- Do you know any of them personally?
- Have you ever been involved in debate on talk pages? [If yes] Could you tell me what happened? (about what, with whom, what the controversy was about, etc.)
- How did you feel?
- How do you 'win' a debate on a talk page?
- After shifting focus from the community to the platform, I decided to interview the Principal System Architect at the Wikimedia Foundation, whose job is to oversee the development of Wikipedia's platform as a sociotechnical system I have focused the interview on design strategy. The interview started with the following set of questions:
 - What's your day-to-day routine?
 - What are you working on right now?
 - How do you make decisions on where work needs to be done?
 - Do you interact with volunteers from the Wikipedia community at all? [If yes] What is the nature of these interactions?

Informal conversations have occurred in the context of contacting volunteers, directly, for interviews. Some of these have been instructive, in that, while not suitable for inclusion in a formal data analysis, have informed some of my conclusions on what matters to Wikipedians.

2.4.3. Documentary sources

The main body of data is constituted by documentary sources. I have already mentioned community pages – policies, essays and guidelines. Additionally, in order to understand the design choices that shape Wikipedia's development, I have examined conversations between Wikipedians geared around technology and design, first-hand accounts of Wikipedia development, and documents outlining design values and requirements. In total, I reviewed and coded 237 documents. Written documents can be grouped, by type, in 5 broad categories: Wikipedia policy, guidelines and essays; community pages; talk pages; documents about Wikipedia; context material; technical documentation.

As explained in chapter 1, Wikipedia policy – comprising policies proper, as well as guidelines – is a crystallisation of practice that has occurred over time, and it has some firm tenets, that are non-negotiable: this means policy pages describing them can't be edited by anyone other than the remaining founder, Jimmy Wales. Essays, on the other hand, are expressions of points of view within the community. While anyone can write an essay, there is a way to gauge how influential they are, by looking at the number of pages that link to them (this number is available for all pages, by clicking on “What links here”, under “Tools”, in the side bar, and then on “Link count”). As a rough measure of scale, the link count for Wikipedia's homepage is currently 878,480; for the main policy, Neutral Point of View, it's 1,581,397; the essay “Wikipedia:Please do not bite the newcomers” has 798,526, while the essay “Wikipedia:Don't bite the researchers” only has 19³. There is also a list of Wikipedia essays which is not exhaustive – those included are likely to be the most popular. Finally, the List of essays page states that there is a filter on essays: duplicated essays are discouraged, and some are deemed too narrow to exist as essays. This process of selection gives them some kind of community support.

Community pages include: category pages, pages aimed at training newcomers, technical literature, user profiles, obituaries, forum discussions. These pages are very useful for information about values that Wikipedians hold dear, as well as common ways of thinking about their editing practice and Wikipedia as a whole. Talk pages are discussion pages related to every Wikipedia page, and their function changes depending on the page they are related to: in the case of articles, talk pages are the venue where discussion on article edits occurs; in the case of user pages, they function more like inboxes – if you want to communicate with someone, write on their talk page. The depth at which I have engaged with talk pages changes depending on what talk pages are involved (more on this further down this section). Under ‘documents about Wikipedia’ I subsume content that is mostly not on Wikipedia, but talks about Wikipedia: that includes Diff (the Wikipedia community blog) posts, the Wikipedia page about Wikipedia, pages from Meta, Wikimedia Foundation documents, writings by Magnus Manske, who designed a large part of Wikipedia's initial platform, and Larry Sanger, who co-founded Wikipedia, as well as press coverage and commentary about Wikipedia.

Among the vast amount of material available on and in Wikipedia, I have selected documents in two waves. At first, I followed the path of a newcomer (which I was), trying to understand how to edit a page and what the rules are. Hyperlinks between pages have been a major driving force in terms of direction, as well as my own curiosity and need for clarification. This operation has allowed me to accrue sources on Wikipedia's culture, and an understanding of the space I was moving in. A second wave of document selection was guided by my research priorities; I selected documents that mapped onto three main subject areas: behaviour, knowledge, and technology.

The first set of documents – conversations – include early email exchanges between Wikipedia's founders, and community pages. An archive of mailing lists containing message exchanges between members of the Wikipedia community are available online at <https://lists.wikimedia.org/>. The archives stretch back to the founding of Wikipedia: they contain conversations among founders about how to structure the project, what kind of features should be developed, and so forth. Messages are signed with names and surnames of individuals. I have initially parsed the content by taking a

³ All numbers obtained from <https://linkcount.toolforge.org>, as of 5 February 2024

cursory look at the content of the conversations, then searched for specific topics that seemed relevant to my questions. As my focus was on aesthetics and design choices, I initially searched generic terms related to Web design: “font”, “appearance”, “markup”. Then, having homed in on specific conversations about skins for the encyclopaedia, I looked for conversations about them, using the names of the skins themselves; the most discussed at the time were “Cologne Blue” and “Standard”. I have then downloaded the most relevant exchanges, and coded them as part of the second thematic analysis, geared around design (see Appendix B). I also searched for “Ayn Rand”, in order to obtain some information as to how much founders were actually interested in her work, and whether she was casually mentioned in conversation. I have also analysed community pages relating to technological development (for full list, see Appendix B). A precious first-hand account of Wikipedia’s development is Larry Sanger’s *The Early History of Nupedia and Wikipedia: A Memoir*, published on the news website Slashdot, a staple of nerd culture (<https://slashdot.org/>). Finally, in order to distil the core values that inspire design and development, I have analysed guidance and strategy documents, including white papers published by the Wikimedia Foundation, product principles, Wikipedia’s Manual of Style, engineering principles, descriptions of specific tools, policies on the usage of tools, and so forth.

As well as looking for specific features, trying to understand how Wikipedia’s platform works has led me to search for technical documentation, regulation and guidance on various coded artefacts, including blog posts, academic papers, Wikipedia and Wikimedia pages. Most of these documents were not included in the thematic analyses, but are referenced throughout the thesis as sources of information, or analysed individually where relevant.

2.5. Data analysis

For this project, I used a mixed methodology, that involved two thematic analyses of textual data, and an analysis of Wikipedia’s website modelled on methods used in platform studies and value sensitive design.

2.5.1. Early work on talk pages

As outlined above, at its inception, the project was geared around the analysis of talk pages. While data from this analysis hasn’t fed into the results presented in this dissertation, early investigations have provided precious context, and can be understood retrospectively as preliminary research. I report here, briefly, how I conducted talk page analysis, for completeness.

The analysis of talk pages can be approached in the same way as analysing correspondence: they are both a conversation, and a permanent record of a conversation. Conversations occurring on talk pages are aimed, mainly, at settling content disputes. Since the role of Wikipedia editors is to collate knowledge from secondary sources, debate ensues is mostly geared around the sourcing and representation of information. The initial analysis of talk pages, when the focus was on rhetorical strategies, was based on argumentation strategies used by Wikipedians in supporting their position. The aim of the analysis was to understand how debate is conducted. Identifying rhetorical strategies would give insight as to how decisions on content are made.

In order to conduct this analysis, I downloaded the entire archive of talk pages for the article

“Pluto” on project management program Atlas.ti, and did an initial round of coding based on the conversations I was reading. Two sets of codes emerged: one concerning rhetorical strategies employed by editors; the other, topics of conversation. As the project took a more sociotechnical character, the coding schemes evolved as well, and ended up being geared around values.

At the same time, given the importance of neutrality, especially as it relates to identity, which has important consequences in terms of the construction of expertise, I then decided to analyse talk pages for article whose content is heavily dependent on the point of view of editors: namely, articles about slurs, which have then been appropriated by the community they were applied to. The pages in question are the articles “Queer” and “N-”⁴. Some of the insights about positionality and its relationship to knowledge-production on Wikipedia stem from this preliminary work.

2.5.2. Thematic analysis

As previously mentioned, I have conducted two thematic analyses: a first one about Wikipedia’s epistemic culture, and a second one specifically geared around design and technology. Of the 122 documents and transcripts which ended up as part of thematic analyses, 99 were analysed for the Epistemic culture thematic analysis (see Appendix A), and 23 for the Infrastructure and design thematic analysis (see Appendix B).

Because of the permeability between social meanings, community culture, and technological development, I used the same coding scheme for both analyses. Naturally, the nature of the data has meant that the distribution of the codes changed depending on the sample. A complete list of codes can be found in Appendix C. The themes emerging from the thematic analyses are also summed up in the respective appendices (A for Epistemic culture and B for Infrastructure and design).

Documents have also served as a way to analyse what kind of values are designed into Wikipedia’s infrastructure. I have described in section 2.3, value-led design approaches have built a methodology to embed, deliberately, values into design. Value sensitive design also allows for a retrospective analysis, aimed at establishing how a certain artefact promotes or hinders human values (Friedman & Hendry 2019).

As well as analysing artefacts in terms of what values they uphold, it is possible to use an approach of a similar level of granularity as the normative value-led design approach. Designers have values of their own, influenced by the community they belong to, and bring those values to their work, by deciding what problems are worth pursuing, imagining the users in a specific manner (Ivory, 2013, 2013; Norman, 2013), and choosing what values are worth considering when creating a product, based on their own ethics and worldview (Costanza-Chock, 2020; Norman, 2013). Thus, I find it helpful to use the value-led design method as an interpretive structure, assuming that, even when a process of value embedding is not explicitly and deliberately followed, a similar path is followed, mostly unwittingly, by designers.

I have combined the spirit of the retrospective analysis with the granularity of their normative accounts, trying to reconstruct, through documents, how some of the key design choices that have shaped Wikipedia have taken place, and how they relate to the ideological and cultural commitments

⁴ Following common practice, as a white person, I have not written out the N-word in my project. However, Wikipedia article uses the word in its full extent.

of its designers, broadly construed as the founders, and the larger community of editors. I have used value-led design methodology as an interpretive frame to analyse the way Wikipedia is built. While Wikipedia is not explicitly designed and maintained based on a value-led design approach, the process by which Wikipedia is designed does present certain features that can be described in terms of value-led design concepts. Conceptualisation, in a similar manner to how it is described in value-led design literature, occurs through articulation of principles that are meant to inspire the design of the platform. Such articulation occurs in a fragmented, but recognisable manner, through discussions about features that should be implemented, design principles listed in technical literature, and strategy documents published by Wikipedia's editor, the Wikimedia Foundation. Discussions, technical and strategy documents explicitly invoke principles that are important to the founders, the community, or to realise Wikipedia's mission, and are used to generate specific interventions on Wikipedia's platform.

Operationalisation (the term I have decided to use, but "specification" would serve equally well) then occurs through the coordinated work of volunteers and Wikimedia Foundation engineers. Due to the decentralised nature of Wikipedia's coordination system, however, sometimes the process is not as linear: coded objects are created to serve the designer's interpretation of Wikipedia's mission, for instance, or to respond to community needs, bringing in new values. Often, those who code for Wikipedia provide explanations of the rationale of their work, by explaining what need they were trying to address, or what kind of ethical principles they had in mind. I will provide extensive examples in chapters 4 and 5.

2.5.3. Platform analysis

Methodologically, platform studies use a kind of hermeneutics: the platform scholar interprets objects and obtains some kind of informed, if partly, necessarily, subjective, insight (Montfort & Bogost, 2009). It is, thus, difficult to define a strict methodology; rather, a series of activities can be listed, which inform the interpretation given as a result of the analysis. In absence of a formalised method, I have set for myself a series of questions meant to guide my analysis, and highlight the cultural influences, deeply-held values, and likely affordances of Wikipedia. The questions are as follows:

1. What is shown? What is hidden?
2. Web pages are designed to display certain parts of the website, and hide others, to different degrees. For instance, the workings of the page, such as conversations between contributors, markup, editing interfaces, can be displayed to various degrees.
3. What style of design is adopted and why?
4. Web design has styles, which have evolved over time, as well as coexisted. Different aesthetics appeal to different audiences, and come from different design principles, which in turn embed ideological assumptions about taste and belonging.

5. What is the artefact similar/dissimilar to? how does it relate to the history of design?
6. Like any designed objects, websites are influenced by other designed objects – websites, but also the wider design landscape: aesthetic movements developed in the context of architecture and furniture design, for instance.
7. What effect might it have on different observers?
8. As mentioned above, styles are deeply connected to audiences. It is important and interesting to try and evaluate, based on the style of a website – in this case, Wikipedia – what the intended audience is.
9. What design principles does it apply or disregard?
10. Design is a highly theorised discipline. Much has been written about the merits of various styles, in terms of usability, aesthetics, ease of design and appeal to specific audiences. Based on what principles seem to regulate the design of specific Web pages, it is possible to deduce how they were meant to be used, whether they were meant to be easily modified, and so forth.

As well as looking at the website from a design point of view, I – like others before me – have paid attention to what kind of affordances it offers. Using a postphenomenological approach, I have also gone further and been able to describe how affordances mediate the production of knowledge and the construction of subjectivity. Values embedded in coded objects can also be evinced from the way the artefacts themselves are designed, by analysing their purpose and functioning in terms of who they serve, for what purpose, how they impact on power relations, what kind of context of use they assume. This kind of approach is often employed in studies in postphenomenology, where the researcher observes the way an object is designed and how it functions in its context of use, in order to draw conclusions on what its role is in the lifeworld of its users (e.g. de Boer *et al.*, 2018; Jensen & Aagaard, 2018; Moerenhout *et al.*, 2020; Smits *et al.*, 2022).

2.5.4. Limitations

The main limitation of this project is that Wikipedia is constitutively mutable: it is made to be changed. At the ground level, Wikipedia's mutability means that, to some extent, any study of Wikipedia is destined to, eventually, become at least in part obsolete – much like ethnographies, or any studies of living organisms. A more sophisticated way of accounting for time is to, on one hand, keep in mind exactly what changes, at what pace; on the other, to fold mutability, and time, into the account itself. On the first count, I have tried to be aware of the half life of the policies, decisions, digital artefacts I have looked at, gauging exactly how far their influence extends. For instance, the strategy documents examined in my analysis were selected based on their validity: while published in 2019 and 2018, they outline a strategy meant to inform Wikipedia's development until 2030, granting them lasting, if necessarily time-limited, relevance.

2.6. Access

Access is not just about the availability of material. It is also about the ability to generate data from it. This tension between availability and readiness-to-hand is particularly important when approaching open source software and the communities around it. In what follows, I list my points of access to Wikipedia, and the interplay between available material and my ability to use it for this study. This theme recurs frequently in scholarship about Wikipedia, and the present study will be no exception.

2.6.1. Background

The Wikimedia Foundation used to have a Research Committee⁵. I tried to get in touch with members of the committee, to seek approval for contacting members of the community, and in order to do so I emailed the mailing list, research-wmf@lists.wikimedia.org, which was associated with the committee. In the response I was given, it was stated that I was not supposed to obtain any formal approval from the foundation, however I was encouraged to submit my project to the project directory (<https://meta.wikimedia.org/wiki/Research:Projects>), in order to allow feedback and advice from the community, before I started any empirical work, as well as emailing the list wiki-research-l-owner@lists.wikimedia.org. I later found out the Research Committee is not active anymore (Wikimedia committees/Research Committee).

I emailed the mailing list, and received the following message in response:

You are not allowed to post to this mailing list, and your message has been automatically rejected. If you think that your messages are being rejected in error, contact the mailing list owner at wiki-research-lowner@lists.wikimedia.org.

I could not see a good reason to press further in this direction. I then uploaded a short description of my project⁶ on the Research:Projects page. I did not receive any feedback, nor was I ever contacted about my project any further.

However, in order to maintain a good relationship with the community while carrying out the project, I sought further guidelines on Wikipedia itself. Wikipedia policy acknowledges the practice of using Wikipedia data and interactions with the Wikipedia community for research purposes. The relevant principle is “Wikipedia is not a laboratory” (Wikipedia: What Wikipedia is not/Wikipedia is not a laboratory), which, at the time of writing, reads as follows:

⁵ <https://meta.wikimedia.org/wiki/Research:Committee>

⁶ https://meta.wikimedia.org/wiki/Research:Sociotechnical_epistemology:_how_do_we_foster_good_practices_in_collective_knowledge-production%3F

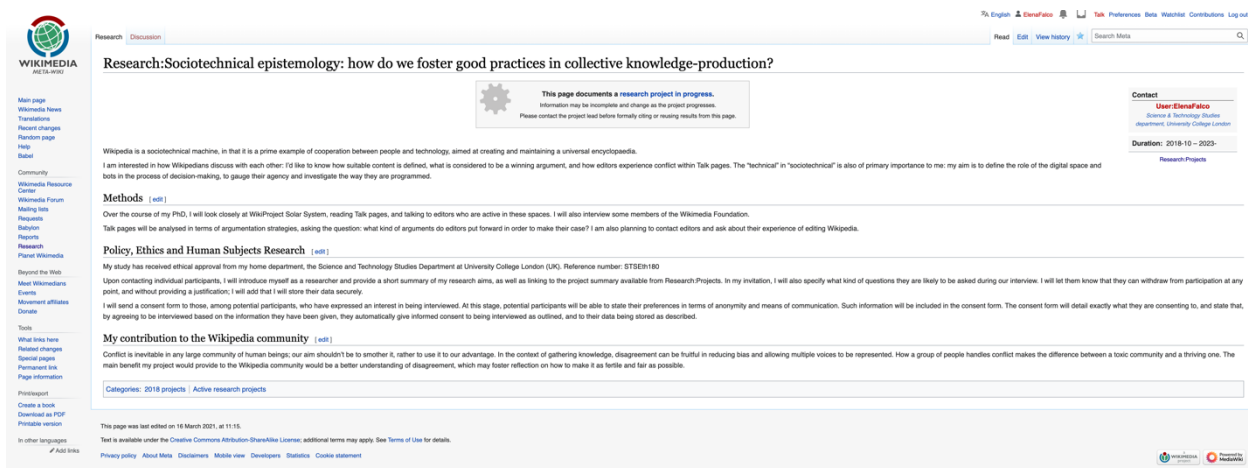


Figure 2.3. My project page on Wikimedia's Meta-Wiki.

Wikipedia is not a laboratory

Research about Wikipedia's content, processes, and the people involved^[7] can provide valuable insights and understanding that benefit public knowledge, scholarship, and the Wikipedia community, but Wikipedia is not a public laboratory. Research that analyzes articles, talk pages, or other content on Wikipedia is not typically controversial, since all of Wikipedia is open and freely usable. However, research projects that are disruptive to the community or which negatively affect articles—even temporarily—are not allowed and can result in loss of editing privileges. Before starting a potentially controversial project,^[8] researchers should open discussion at the Village pump to ensure it will not interfere with Wikipedia's mission. Regardless of the type of project, researchers are advised to be as transparent as possible on their user pages, disclosing information such as institutional connections and intentions.^[9]

Some editors explicitly request to not be subjects in research and experiments. Please respect the wish of editors to opt-out of research.

(Wikipedia: What Wikipedia is not/Wikipedia is not a laboratory)

While I did seek to connect with Wikipedians before starting empirical work, I did not open a discussion on the Village pump, for two reasons: I did not envisage my work to be disruptive, and I don't have an allegiance to Wikipedia's mission. In my view, social researchers should be respectful of their participants, and owe them a duty of care as individuals, but are free to critique institutions, and institutional practices, provided their arguments are rooted in rigorous scholarship. While the boundary between community and institution are blurry on Wikipedia, discussing my project with the aim of not interfering with Wikipedia's mission would, potentially, have led to modifications of the project design made in accordance with Wikipedia's priorities: or, in other words, would have constituted a conflict of interest. Conversely, had I not followed the advice given, I would have exposed myself to unnecessary criticism, and potentially even harassment⁷. As further interactions

⁷ Multiple sources, internal to Wikipedia, to whom I have promised anonymity, have mentioned targeted harassment within the community, including in-person stalking. Some cases have become widely known, others haven't. Unlikely as it is for me to be a victim, I decided that, on balance, it was worth keeping a low profile.

with Wikipedians have shown (see details in section 2.6.2), I was correct not to engage. I will more extensively detail my position on this matter in section 2.6.3.

2.6.2. Recruitment of participants

In order to facilitate access to participants, and be as respectful as possible in conducting my research, I have sought other resources on the matter, and followed the advice given in related essays, namely “Wikipedia: Research” (Wikipedia: Research/Advice for researchers) and “Wikipedia: Ethically researching Wikipedia” (Wikipedia: Ethically researching Wikipedia).

My approach to contacting participants differed based on their role within Wikipedia and the Wikimedia Foundation. I have contacted technical and non-technical staff at the Wikimedia Foundation, and volunteers. In the case of non-technical Wikimedia Foundation staff, informal introductions started the process: two colleagues who used to be active participants respectively in the Argentinian, and UK Wikipedia communities referred me to one member of staff each; another member of staff approached me on Twitter, where I am transparent about my research and occasionally engage with Wikipedia-related content. It is worth noting that the staff members I ended up talking to are in public-facing roles, such as communication and community liaison. Snowballing and introductions took care of the rest.

The process was different when it came to talking to Wikimedia staff in the product and technology teams. Informed by my understanding of office etiquette, rather than cold emailing a large number of people in said teams, I decided to email a senior member of staff, explaining my work, linking to information about my research project, and asking whether it would be appropriate for them to introduce me to staff who might be available for interviews, or if they would be available to participate themselves. The first response was somewhat dismissive: despite having specified that I wanted to talk to those who work for the Foundation, I was invited to email the mailing list of volunteers who contribute technical work to Wikipedia, wikitech-l@lists.wikimedia.org, or talk to the Research Team during their office hours. Only after thanking them and reiterating that, for my purposes, talking to engineers would be best, I was referred to a member of staff, and told another member of staff may also be available. I eventually managed to speak to the member of staff explicitly mentioned in the email, but wasn't put in touch with the other member of staff, despite explicitly asking for them to do so. As speaking to volunteers may also have been beneficial to my study, I have also emailed the technology mailing list. I have not received any answers.

Finally, I have approached editors through their talk pages. Their reaction to invitations informs the broader question of access. Out of 15 Wikipedia editors I contacted directly, without introductions, only 2 ended up talking to me at length, and filling in a consent form. Of the remaining 13, 3 replied to my invitations, but eventually declined to speak at length. I sent 2 further invitations as a result of snowballing; I didn't receive a response from either.

Those editors who replied to my invitation, but eventually declined to participate, showed a similar attitude: they questioned my methodology, or my decision to approach them, or both. One of the two editors I eventually interviewed initially replied in quite a hostile way to my invitation, questioning the format of the consent form and the objectivity of my research. Only when I suggested that we part ways, they apologised for their behaviour and we had a fruitful conversation.

Finally, I approached three people involved in the foundational work on Wikipedia, during its

inception in the early 2000s. I did not obtain a response from any of them.

2.6.3. The politics of access

My experience in attempting to reach Wikipedians highlights a tension in Wikipedia's culture, between openness and access. Formally, Wikipedia is as open as it can be, as a platform, a community, and an object of study: anyone can edit its contents; people who occupy positions of power are expected to be approachable (one of the participants told me how there is an expectation that even the CEO of Wikimedia Foundation reply to messages on her talk page); all data, except users' personal data, is available to download and reuse.

However, this openness is only nominal. Editing content is possible only for a subset of people, with the right material conditions, culture, language, and understanding of the (sprawling and informal) rules of Wikipedia. Talking to Wikipedians is not easy – the mere fact that a venue for communication (talk pages) exists doesn't guarantee a dialogue; access privileges (in the technical sense) are not equally distributed. Finally, as described above and in chapter 1, most ethnographic work about Wikipedia has been carried out by members of the community.

I should note that I am not the only person who received resistance, as a researcher, from the Wikipedia community, as indicated by the existence of – admittedly low on link numbers, at 15 and 19 respectively – essays such as “Wikipedia: Don't bite the researchers”, and “Wikipedia: What are these researchers doing in my Wikipedia?”. The former essay's title is a play on “Don't bite the newcomers”, a common plea to established editors to be kind to new editors (Wikipedia: Don't bite the newcomers) – and, meta-textually, hints at how the opposite attitude is the norm. These two essays try to explain why Wikipedia is interesting for researchers, and to show how academic research may be useful to Wikipedia. Given the small size of the “Researching Wikipedia” category (9 pages), and the relatively significant number of links to each (20 link to the former, 11 to the latter, while “Wikipedia:Research”, which is the main port of call for researchers on Wikipedia, has 203 links), they are relatively important essays on this specific topic.

On to my personal experience: the first thing that jumped up at me was that, among those who replied to me, resistance was never expressed with a simple refusal, which may be attributed to any sort of reasons such as lack of time, energy, unwillingness to talk to strangers. Refusals were always qualified, with, as I mentioned above, specific criticisms of my approach. In other words, objections were raised in epistemological language. It may be that the people in question thought of me not as an epistemic peer, with their own understanding of the world and epistemic objectives, but as an epistemic inferior; someone they were asked to teach to. Elsewhere in my meanderings on Wikipedia, I have encountered a similar attitude towards academia in general, highlighting an alleged lack of objectivity, perhaps compounded by a perceived sense of superiority from academics themselves. Withdrawing participation, then, may be both a way to assert one's epistemic status, and to undermine the project of someone conceived in adversarial terms.

Another interesting aspect of how refusal was framed was the use of technical language: one of the participants raised the issue of my using proprietary software (Word, developed by Microsoft) for the consent form. This kind of resistance both signposts belonging to open source culture (and, consequently, my not belonging to it) and hides an emotional response. Further, they questioned my assumptions: the fact that I had assumed the other person had Word was a marker of my bad

character. While clearly the person in question harboured some kind of anger towards me (later on, they apologised for their attitude – as one does when one has a misplaced affective response, not when one has technical issues), they couldn't express it as such: they turned it into a technical problem.

Diffidence itself deserves some attention here as well. Even once I was talking to participants in amicable ways, they often qualified what they were saying, pre-empted objections, raised issues with Wikipedia that I had not raised in the first place, evoking criticism of Wikipedia found in public discourse. Wikipedia has, indeed, been widely criticised on grounds of bias, and representational injustice (Bear & Collier, 2016; Ford & Wajcman, 2017; Hill & Shaw, 2013). One of the volunteers I interviewed talked at length, unprompted, about bias on Wikipedia; the same participant held beliefs about what is often called “liberal bias” in academia, that is the idea that objective truth is cast aside for ideological reasons, e.g. by hampering research on gender difference. Given my positionality, as a feminist woman in the social sciences, with progressive views, and an online presence that reflects them, it is possible that participants may have worried about my intentions; they may have worried, in other words, that my aim was to criticise Wikipedia, and that they may not be able to convince me not to. Hence the diffidence.

In conclusion, in a community where relationships and shared ideological common ground are highly valued, it is unsurprising that most people approached were not keen to be interviewed. In my use of interviews, I have tried to keep in mind the influence of selection bias: those who have been generous with their time are probably on the most progressive and self-reflective end of the spectrum, within the Wikipedia community. Approaching (and being rejected by) Wikipedians also provided clues about their attitude towards those who are not members of the community, and academia – both aspects of Wikipedia culture will be present in my analysis.

2.6.4. Access to material sources

As previously stated, most data on Wikipedia is available and easy to access. The difficulty is at the level of interpretation: looking at talk pages is not sufficient to fully understand what's going on. Firstly, talk pages are edited themselves. They are the result of tidying up and reshuffling of conversations into a narrative. Secondly, the content of the conversations themselves is fully intelligible only with a thorough understanding of policy and the basics of how Wikipedia works from a technical point of view (see Ford & Geiger, 2012). For instance, the acrimony often found in talk pages can be understood by keeping in mind that some of the people debating are debating their own edits – that is, their own work – and only after they have been reverted a few times – that is, after someone they don't know has implicitly disagreed with them multiple times. For the technical side of things, the importance of previous knowledge is even more pronounced: it's very easy to see Wikipedia's back end, and how it all goes together, through a set of coordination tools of the same kind used by programmers everywhere. But in order to grasp what is going on, some technical knowledge of coding is required.

Acquiring the first skillset – knowledge of Wikipedia as an environment – took some time and effort, as described in section 2.4.1. As for the technical literature, I had some previous, basic coding knowledge (I have taken an SQL course in the past, done some data analysis, and worked in roles that required me to write HTML and markup used in the production of ebooks); as well as an

understanding of the workflow and terminology of software development. During my PhD, I have learnt the basics of Python and JavaScript. I am not an engineer, but I speak the language.

The following may seem frivolous, but it's important: I was a teenager at a time (the early 2000s), when, while social media was in its infancy (Facebook became available to everyone in 2006, the year Twitter was founded), socialising online was already widespread among my peers. In order to open and maintain a MySpace profile (which I did) and participate to online forums (which I did occasionally), one had to learn a bit of HTML; familiarity with markup came with the territory. This is not true of Internet natives who started their online socialising on Facebook, Twitter, TikTok, where the source code is completely opaque. I also socialised with people within the open source and hacking communities, and have thus become fluent in the lingo and accustomed to a certain way of talking about software. As I highlight in chapter 4, this kind of cultural background helps Wikipedians do their job, and arguably its absence hinders the ability of newcomers to join. In my case, this cultural background made it easier for me to carry out research.

2.6.5. Ethics

As this study partly relies on interaction with human participants, I have sought and obtained ethical approval from my home department. Beyond institutional requirements, and in line with best practices in social research, I have reflected on how to approach participants and documents in an ethically sound manner. Specifically, I have reflected on two ethical issues: one concerning power, and one concerning privacy.

Much social research informed by a commitment to justice operates under the assumption that there is a power asymmetry between researcher and participants, the former being in a position of default dominance over the latter. Researchers have power, within this framework, because they extract data from participants, control the conditions of the interaction, and the ultimate representation of the community being studied. A commitment to equity, then, induces many researchers to reduce the asymmetry by, for instance, compensating participants for their time, co-producing the design of the study, asking for, and incorporating, feedback on the results of their study before publication.

As a researcher committed to equity, I have evaluated whether any such measures would be called for in my engaging with Wikipedians, and employees of the Wikimedia Foundation. The decisive question, in this sense, is the following: is there a power differential between me, as a researcher, and the Wikipedia community? The answer is not straightforward. The Wikipedia community is composed of volunteers, and employees of the Wikimedia Foundation. My work, however, doesn't relate the experiences of individual Wikipedians, be it volunteers or employees – it concerns the organisation as a whole. Hence the main object of study – and, consequently, the object of critique in this dissertation and any further publications – is Wikipedia as a project, and the Wikimedia Foundation as an organisation. There are of course links between the organisations and the people who run them, but the difference is important. I, as a researcher, am clearly less powerful than Wikipedia, one of the most read and influential websites in the world, or the Wikimedia Foundation, who publishes Wikipedia.

I, of course, have a duty towards individuals to abide by the boundaries set when establishing consent, and to represent their views fairly. Complying to institutional ethical standards, as set in my

ethics statement, addresses the former duty. As for the latter, I have offered my participants the opportunity to read drafts of my publication, with the intention of incorporating any feedback regarding fair representation.

Clarifying this point is particularly important because Wikipedians often frame themselves, both individually and as a community, as underdogs, or, as one of the ethnographies written by Wikipedians puts it in the title, “a bunch of nobodies”. This is not true, and has never been true: the initial financial impetus for Wikipedia came from the profits of a successful Web portal, called Bomis, co-founded by Jimmy Wales, off the back of a career in finance. Over time, as noted above, Wikipedia became one of the most read websites in the world; its cultural influence is enormous. Wikipedia editors are also a demographically privileged community: “if there is a typical Wikipedia editor, he has a college degree, is 30- years-old, is computer savvy but not necessarily a programmer, doesn’t actually spend much time playing games, and lives in US or Europe” (Wikimedia Foundation, 2011).

The second, related, point, is privacy. It is common in netnography – Web ethnography – to use as data conversations that occur in spaces like social media, blogs, and forum. The debate over how to ethically use this kind of data is far from settled. I will home on one specific point here: the grey area between public and private content. Content found online, even when easily available, is not necessarily meant for public consumption: this is true, for instance, of niche forums, where, even in absence of registration, members reasonably expect their posts to be read exclusively by other members of the same community. A rule of thumb often used by researchers is to ask oneself what the intended audience of a piece of media was at the time of writing: if, for instance, someone wrote a post on their MySpace page when they had 25 connections, it makes sense to assume that the content of the post wasn’t meant for widespread consumption. Hence, consent would need to be sought in order to include the content in one’s dataset.

When it comes to the written sources I have used, the principle of assumed audience seems to me to be relevant to three kinds of text: conversations on talk pages, conversations on user talk pages, and email exchanges on mailing lists. When writing on talk pages, editors know that they are contributing to a public forum: the aim of a talk page is to discuss the content of the page with anyone who wishes to edit it, or weigh in. Talk pages are also archived, and occasionally tidied up in threads to that end, because they are a record of how a specific piece of information was shaped by the community. As I argue in chapter 4, they are a form of validation. Talk pages, then, are inherently public.

User talk pages are a different matter: while, structurally, they are identical to article talk pages, they are a special case, in that they are used for direct communication between editors. From a material point of view – the practical ability to read them – they are public. However, if considered as social spaces, they are, in my view, to be considered private: reading an exchange between two people on their respective talk pages is equivalent to eavesdropping into a conversation on the street. Consequently, I have not used user talk data in my analysis; not even conversations I had⁸.

⁸ Different rules would apply to conversations on talk pages of well-known users, as they often function as spaces to air community grievances. Maintaining the metaphor of conversations on the sidewalk, messages on prominent Wikipedians’ user talks, such as, for instance, the CEO of the Wikimedia Foundation, can be likened to questions

Email exchanges on mailing lists are a special case. At the time of writing, one would typically presume that the authors of said emails would only have the restricted audience of their peers in mind. However, other considerations need to come into play. First, the email exchanges were made available, presumably, by some of the same people who participated in them: archived conversations are made available by the Wikimedia Foundation. I happen to have read exchanges that occurred when not many people were involved in Wikipedia. Secondly, the conversations I read were work, not private, conversations, which ended up shaping an influential artefact; their status as a record supplants the significance they had to those who participated at the time. Third, it is likely that part of the reason why those conversations are public is a commitment to transparency, often abided by in open source communities, which entails a sufficient level of comfort with a potentially extended readership. Finally, early conversations occurred between Wikipedia founders whom, as argued above, I do not consider to be vulnerable; my attitude would be more careful if, for instance, I were approaching a more recent mailing list where a subset of volunteers were discussing a specific strain of editing, especially if about a sensitive topic. Since I kept my focus on powerful figures, I believe I have acted in an ethically sound manner, and in keeping with a spirit of scholarly independence.

Conclusion

In summary, the study presented in this dissertation was conducted by using a mixed methodology, developed on the basis of a composite theoretical framework, which integrates postphenomenology, value-led design approaches and critical approaches to the analysis of digital objects (critical code studies and platform studies). Scholarship rooted in each component of the theoretical framework favours specific methodologies: postphenomenology makes ample use of the concepts of affordance and mediation; value-led design approaches, and value sensitive design in particular, use technical analysis, and engagement with designers (broadly construed as the community around the development of a specific artefact); critical code studies and platform studies analyse digital objects as cultural artefacts, by examining their history, context of production, and intended audience. The approaches used here overlap theoretically in terms of a commitment to materiality, situatedness, and relationality, and methodologically in their shared preference for qualitative, observational methods rooted in aesthetics, semiotics, and hermeneutics.

shouted to the speaker at a public event. For present purposes, however, this point is moot: I have not read nor made use of conversations on the user pages of prominent Wikipedians. I include this note for completeness.

3. Wikipedia's aesthetic

I collected from the diverse flowers of sacred scripture and philosophic writings this book [...] and for the sake of your love as if into a single sweet honeycomb. Therefore, in this very book, you ought diligently to seek pleasing food and to refresh your exhausted soul with its honeyed dewdrops, so that, always occupied with the caresses of the Bridegroom and fattened on spiritual delights, you may cheerfully hurry over ephemeral things to possess the things that last forever in happiness and pleasure.

Abbess Herrad of Landsberg, Hortus Deliciarum

3.1. Introduction

In this chapter I will examine how Wikipedia's aesthetic responds to the project's epistemic commitments, strategic priorities and changes in the technological landscape. By "aesthetic", I mean Wikipedia's visual appearance and any stylistic choices related to it, including the way in which Wikipedia is written, and branding guidelines. When I use the term "aesthetics", I refer to parameters and ideals pertaining the experience of a designed object. This chapter is informed by a thematic analysis conducted specifically on documents related to Wikipedia's design and branding – which include white papers, branding guidelines, the Manual of Style, and mailing list conversations about design – and conversations with and among Wikipedia's about design.

I suggest that Wikipedia's interface responds to worries around relevance, by aiming to expand readership, retain readers, and turn them into contributors. I then show how a tension between abstraction and situatedness defines the way content is displayed. Subsequently, I explore how the contours of the Wikipedia community intersect with design choices, through the construct of the imagined user. I conclude by placing Wikipedia's aesthetic in the same lineage as modernist design, exploring the historical and ideological ties between modernism, Ayn Rand, the free and open source software movement, and Wikipedia.

The overarching theme of Wikipedia's aesthetic, in my view, is an attitude that Sherry Turkle defined as "antisensuality": a disregard for the body, detected by Turkle in nerd communities, which reverberates through Wikipedia as the concealment of the bodies of editors, the downplaying of materiality, and the performance of abstract ideals of neutrality and universality.

3.2. Interface design and Wikipedia's sustainability

In this section, I will argue that the design of Wikipedia's interface responds to a key concern about the project's survival: maintaining a readership, which can, in turn, be converted into a pool of editors. In order to address these concerns, the interface aesthetic is designed to highlight the relevance of Wikipedia's content, keep readers on the site, and invite them to edit. Wikipedia is

presented as relevant, in the sense that it contains information that is important to potential readers, useful to make sense of their lives, and up to date. Secondly, Wikipedia's design welcomes and stimulates curiosity, by offering opportunities to the reader to venture beyond the piece of information they were initially seeking. Third, Wikipedia presents itself as malleable and unfinished – as something that invites, and requires, intervention.

Worries around the survival and sustainability of Wikipedia's project are expressed in a white paper called *Scale*, published by the Wikimedia Foundation in 2019, and explicitly related to decreasing levels of attention:

it will be critical to design for resilience: the ability to engender sustainable growth and fend off threats. For example, it will be necessary to define countermeasures against external threats such as censorship, misinformation, climate and policy related threats, as well as attacks on security or privacy by state actors. It will also be necessary to anticipate and countermand threats that could undermine the projects from within: communities or affiliates turning against one another, communities turning against themselves and communities turning against the Foundation. And finally, perhaps the most critical existential threat is relevance; what barriers to entry can be erected to prevent loss of mind share? What pre-emptive measures must be taken to guarantee mind share as new communities come online? [...] As the incumbent nonprofit Internet presence defending a neutral point of view and access for all, it is critical that Wikimedia maintain and strengthen itself to preserve a future with truly free knowledge.

(MediaWiki: Scale [white paper], p. 15)

■

This quote shows the level of concern felt within the Foundation: perceived threats come from outside and inside the community, but the most felt is relevance: the idea that people's attention ("mind share") might become scarce, hence reducing access and threatening the survival of the project. Conversely, since Wikipedia's only way to increase the number of contributors is to invite readers to contribute, fewer readers mean fewer opportunities to recruit new editors. Concerns about the number of editors have been present for a long time within the Wikipedia community (see Appendix A), and still are. Such concerns often inform decision-making: the influence of a given practice or initiative on the uptake of new editors constitutes a stock argument in conversations among Wikipedians.

I suggest that worries around sustainability underlie the semiotics of Wikipedia's interface, which seems designed to attract and retain readership, as well as turning readers into contributors, through accessibility and engagement.

3.2.1. Accessibility

For a website to attract readers, the ground-level requirement is for it to be readable. Accessibility is considered a fundamental, and uncontroversial, value in Web design practice; at the time of Wikipedia's inception, accessibility was arguably the main ethical concern for Web designers. Partly

because of this context, and partly because of Wikipedia's need to grow a large editor base, and, consequently, to be able to be usable by as large a group of people as possible, accessibility has informed the design of Wikipedia's interface in various ways, which I will describe over the course of this chapter.

When Web design first started to be discussed in ethical terms, accessibility was one of the main concerns (Kennedy, 2012). Accessibility was defined specifically with respect to disabilities that might make it difficult to interact with websites, and aesthetic choices were the primary topic of discussion: a minimal appearance is essential for a website to be accessible (*ibid.*). Implementing accessibility can translate to several different technical requirements. For instance, high contrast between text and background, with minimal distractions, is deemed to help people with minor visual impairments (who can see, but struggle to tell certain objects apart). Websites can be designed to allow people to use screen readers: this means using actual text as opposed to images of text, and alternative text to describe images.

In line with this way of thinking, Wikipedia's Manual of Style prescribes writing in a way that is compatible with screen readers:

Avoid article text referring to images as being to the left, right, above, or below, because image placement varies with platform (especially mobile platforms) and screen size, and is meaningless to people using screen readers; instead, use captions to identify images.

(Wikipedia, Manual of Style)

By avoiding language pointing to the placement of images on a page, Wikipedia is intelligible to those using screen readers, who hear the page as opposed to seeing it. Similarly, colour coding shouldn't be used to express meaning:

When making user-made diagrams or similar images, try not to use color alone to convey information, as it is inaccessible in many situations.

(Wikipedia, Image use policy)

Text should be typed instead of placed in images:

Generally speaking, you should not contribute images consisting solely of formatted or unformatted text, tables, or mathematical formulas. In most cases these can instead be typed directly into an article in wiki markup (possibly using MediaWiki's special syntax for tables, math). This will make the information easier to edit, as well as make it accessible to users of screen readers and text-based browsers

(Wikipedia, Image use policy)

Considerations around accessibility are used to constrain the design of coded artefacts as well.

Bots, for instance, are meant to respond to accessibility worries. When designing and deploying bots, Wikipedians need to be mindful of the impact the bot’s activity will have on the page. Often, restrictions on bots design factor in assistive technology, as shown in these quotes from Wikipedia’s Bots policy:

the output text or HTML in ways that make *a difference* to the audio or visual rendering of a page in Web browsers, screen readers, when printed, in PDFs, or when accessed through other forms of assistive technology”

Changing HTML entities to Unicode characters whenever the Unicode character might be difficult to identify visually in edit-mode, per the Manual of Style.

(*Wikipedia: Bots policy*)

Accessibility permeates Wikipedia’s design and design practices. From an aesthetic point of view, it translates into a plain appearance and sparse use of images, that are then turned into propositional content for screen readers – images are chosen for their informational, rather than visually appealing, qualities. Over the course of this chapter, I will show how central a minimalist appearance is to Wikipedia, for reasons related to ethical values (here: accessibility), community practices, and the broader cultural landscape Wikipedia was designed within. I will also show how the tension between minimalism and broad appeal underlies much of Wikipedia’s aesthetic conundrums. I will start, in the next section, by describing how Wikipedia’s interface attempts to draw, and retain, readers and editors.

3.2.2. Engagement

In order to attract and retain readers, Wikipedia’s content needs to be presented as relevant to their interests, and invite a longer engagement with content than a quick check. A good illustration of how design elements can be conducive to these goals is Wikipedia’s Main Page. The default Wikipedia home page is shown in Fig. 3.1. It has multiple sections. In the middle is highlighted content: lightly coloured rectangles enclose snippets of content. In green, to the left, are a featured article, and a “Did you know...” section with trivia; in light blue, to the right, are the “In the news” section, with a selection of news items, and the “On this day” section, which highlights some events that happened on the same Gregorian calendar day, in the past. The lilac section at the bottom contains the day’s featured picture, with a brief description of its relevance.

The blue boxes on Wikipedia’s interface give the reader a sense of time. The top box highlights the present (“In the news”), showing how Wikipedia can help making sense of current affairs: each news item contains links to pages about the news’ protagonists. The bottom box (“On this day”) is focused on the past: there is a logic of “learning from history”, to an extent – but also an element of arbitrariness. The calendar is a construction, and just because something happened on this day, in 1789, it doesn’t mean it’s particularly relevant today. “On this day” provides an excuse to look at history; an artificial sense of salience which points towards a more profound way in which history is seen as important to understand the present. The green boxes, curated by editors and

Welcome to Wikipedia,

the free encyclopedia that anyone can edit.
6,789,596 articles in English

From today's featured article



Female polar bear

The **polar bear** is a large bear native to the Arctic and nearby areas. Closely related to the brown bear, the polar bear is the largest extant species of bear and land carnivore, with adult males weighing 300 to 800 kg (700 to 1,800 lb). It has white or yellowish fur with black skin and a thick layer of fat. Polar bears live both on land and on sea ice, and usually live solitarily. They mainly prey on seals, especially ringed seals. Male bears guard females during the breeding season and defend them from rivals. Mothers give birth to cubs in maternity dens during the winter. The International Union for Conservation of Nature considers polar bears a vulnerable species. Their biggest threat is climate change as global warming has led to a decline in sea ice in the Arctic. They have been hunted for their coats, meat and other items. They have been kept in captivity and have played important roles in culture. (Full article...)

Recently featured: "Last Gasp" (Inside No. 9) · Benty Grange hanging bowl · Nestor Makhno

Archive · By email · More featured articles · About

Did you know ...

- ... that **Frankie Connolly** (pictured) signed to *Xenomania* and took her first singing lessons – in that order?
- ... that almost all members of a flock of **tufted jays** work together to build a nest?
- ... that **Sarah McCreaner** imitates objects under *hydraulic presses* through dance?
- ... that former US president Theodore Roosevelt was shot in front of the **Glipatrick Hotel** in 1912?
- ... that Olympic **hammer thrower Tamer Balci** was later cast in a movie as **Tarzan**?
- ... that **TV stations in Nashville** and in **Memphis, Tennessee**, both on channel 30, lost their **Fox** affiliations in 1990?
- ... that the mild climate in the remote mountain village of **Wakan** in Oman allows the growth of fruits such as **pomegranates** in an otherwise hot and dry country?
- ... that football player **Brad Kragthorpe** played in an exhibition game named after his grandfather?



Frankie Connolly

Archive · Start a new article · Nominate an article

In the news

- Following the **general election**, **Feleti Teo** (pictured) is appointed **Prime Minister of Tuvalu**.
- The ***Odysseus*** robotic lander of the **IM-1** mission performs the first commercial soft landing on the **Moon**.
- At the **British Academy Film Awards**, **Oppenheimer** wins **Best Film** and six other awards.
- Russian opposition leader **Alexei Navalny** dies in a **corrective labor colony** near **Kharp**, at the age of 47.



Feleti Teo

Ongoing: **Israel–Hamas war** · **Myanmar civil war** · **Red Sea crisis** · **Russian invasion of Ukraine** (timeline)
Recent deaths: **Hydeia Broadbent** · **Edith Ceccarelli** · **Andreas Brehme** · **Michael O'Regan** · **Bobbie Wygant** · **Jack Biddle**

Nominate an article

On this day

February 27: Feast day of **Saint Gregory of Narek** (Catholicism)

- 1776 – American Revolutionary War: A Patriot victory at the **Battle of Moore's Creek Bridge** resulted in the arrests of 850 Loyalists over the following days.
- 1814 – **Peninsular War**: In the south of France, Spanish, British and Portuguese soldiers under the command of **Arthur Wellesley** defeated French soldiers in the **Battle of Orthez**, causing the French to retreat east.
- 1972 – *The Sunday People* revealed that **James Humphreys**, a notorious pornographer, had bribed a senior London police officer with a holiday to Cyprus and Beirut.
- 1988 – The Armenian community of **Sumgait** in Azerbaijan was the target of a **violent pogrom**.
- 1996 – The multimedia franchise ***Pokémon*** was launched with the release of the video games ***Pocket Monsters Red and Green***.




Memorial to the victims of the Sumgait pogrom

Robert of Melun (d. 1167) · **Alice Hamilton** (b. 1869) · **Ganesh Vasudev Mavalankar** (d. 1956) · **Leah Poulton** (b. 1984)

More anniversaries: February 26 · February 27 · February 28

Archive · By email · List of days of the year

Today's featured picture



The **Malaspina Glacier** is the largest **piedmont glacier** in the world, situated at the head of the ***Alaska Panhandle*** in the United States. It is about 65 kilometers (40 mi) wide and 45 kilometers (28 mi) long, with an area of 3,900 km² (1,500 sq mi). The glacier is up to 600 meters (2,000 ft) thick in places, with the elevation of its bottom being estimated to be up to 300 meters (980 ft) below sea level. This three-dimensional rendered panoramic map shows the Malaspina Glacier and surrounding area in 2021 – the glacier is at the front and center, with its concentric rings of ice, rubble and meltwater visible. The map was created using a digital elevation model with height samples every 22.5 metres (74 ft) based on United States and Canadian government data. The natural colors represent water, snow and ice, barren land, low vegetation, broadleaf forest, needleleaf forest, and wetland.

Panoramic map credit: Tom Patterson

Recently featured: *The Grands Boulevards* · *Anax ephippiger* · *Lychee*

Archive · More featured pictures

Figure 3.1 Wikipedia's main page.

displaying articles of their own choosing, encourage exploration by showcasing the breadth of information that can be found on Wikipedia.

The affordances and expressivity of the Main Page are enabled by hyperlinks. Without them, the homepage would be unable to perform any of the functions outlined above, and probably look very different. The news items, for instance, would be mildly informative by themselves, but lose meaning in context: why put news on the frontpage of an encyclopaedia? They only make sense if linked to the contents of the encyclopaedia itself. The same can be said about the “On this day” section. The featured article would need to be shown in its entirety, in order to make sense. Only the self-contained elements would hold up: the short facts to the left, and the picture at the bottom, would still count as entertaining.

There is a co-dependency between Wikipedia and hypertext: they play to each other’s strengths. Hypertext doesn’t work well with every kind of content: the promise of hypertext generated an initial enthusiasm, and consequent experimentation with its narrative possibilities (see, e.g., Joyce & Tringham, 2007). However, only some stuck. There are kinds of cultural products that lend themselves to shifting power to their audience, while others don’t: hypertextual novels have very limited appeal, for example; videogames are very popular; we are currently in an experimental phase for video content that allows the audience to choose (like Netflix’s *Bandersnatch*), but this area is still minuscule, compared to vast majority of storytelling produced in the world.

Wikipedia content is suitable to be consumed in hypertext format, because, differently from novels, it is not sequential: it’s modular, made of little titbits that can be understood without being consumed one after the other. Modular content can be reconfigured and realigned based on the reader’s interests: one doesn’t read one article after the other in a specific order – rather, one’s interests and curiosity lead the way. The route a reader takes through Wikipedia is a map of the reader’s curiosity. The importance of this fact will become apparent in the context of the tension between abstraction and situatedness (section 3.3).

To an extent, articles are not linear either: while of course they are written sequentially, they are themselves scattered with links. Links can lead directly to a certain section of the article, and openings to veer off in different directions, away from the narrative of the article itself are offered throughout. The reader can find themselves mid-sentence, encounter a hyperlinked word, marked in blue, click on it, and find themselves in a completely different place, reading an article about the word itself.

The symbiosis between hyperlinks and Wikipedia allows for two widespread cultural practices: the social fact-check and the rabbit hole. Wikipedia is often used in social contexts, to confirm the accuracy of a statement, or find out more information about a topic of conversation. Encyclopaedias have long had a social function: part of their appeal has always consisted in enhancing social life by allowing compact knowledge-titbits to travel through the middle classes. Propositional knowledge, expressed through written text, lends itself to this usage which has changed over time, but has maintained a relationship between social interaction and the unit of the “fact” as a titbit that can be easily shared with others. In the case of Wikipedia, hyperlinks, and the larger infrastructure built upon them – the Web, and search engines – allow for single units of information to be found quickly, and spent in a social environment.

The Wikipedia rabbit hole consists in the practice of starting from one place, and following one’s

interests, clicking on a sequence of hyperlinks in order to explore the topic further, or jump from related topic to related topic. The importance of the Wikipedia rabbit hole is attested by its acknowledgment in marketing copy: one of the introductory screens for the Wikipedia app, published by the Wikimedia Foundation, explicitly invites users to delve into one (Fig. 3.2).

The description depicted in Fig. 3.2 also shows how one of the aims of Wikipedia's design is to encourage exploration, and connects exploration to the features I described above (here only "On this day" is mentioned, but it's reasonable to believe that other similar devices fulfil the same function). Papson, Goldman and Kersey (2004) argue that hypertext (which I am here reducing to hyperlinks, as the essence of hypertext) shifts the power from the writer to the reader: rather than following the text in a linear manner, the reader can choose their path – explore. Theodor H. Nelson, who invented hypertext, explained his rationale in his manifesto *Dream Machines*, in which he



New ways to explore

Dive down the Wikipedia rabbit hole with a constantly updating Explore feed.

Customize the feed to your interests – whether it's learning about historical events **On this day**, or rolling the dice with **Random**.

Figure 3.2. Introductory screen to the Wikipedia app, published by the Wikimedia Foundation.

claims that, since we don't think linearly, we shouldn't have to write linearly (1974). The role he initially envisaged for hypertext was as a way of writing in parallel, drawing connections between multiple documents through hyperlinks. *Dream Machines* is written in a manner that follows the logic of the hypertext: partly typed, partly hand-drawn, it does not follow a linear style – titles are sometimes in the middle of the page, multiple thoughts are followed in parallel in different parts of the page, and arrows are present that link separate parts of the document.

Hyperlinks afford a level of autonomy to readers, who become, by virtue of designing their own experience, producers of knowledge. Hyperlinks are not just a way of retrieving knowledge – they

are also a way of creating it:

Uneducated people typically think of education as the learning of a lot of facts and skills. While facts and skills certainly have their merits, “higher education” is also largely concerned with tying ideas together, and especially alternative structures of such tying-together: with showing you the vast uncertainties of things.

(Nelson, 1974, p. 48)

The memex, ancestor of hypertext, presented a similar affordance: it allowed users to find their way through a large corpus of information, and, crucially, to record their path, which could then be reused or passed on to others (O’Sullivan, 2011). Making one’s own way through Wikipedia is the first step to creating content.

For users to feel inclined to contribute, however, two things need to happen: Wikipedia needs to look malleable, like something that affords intervention; and it needs to look like it needs intervention. Users, in other words, need to feel like they can, and should, edit Wikipedia. Wikipedia’s malleability is signalled by literal calls to action such as the “edit” buttons, which addresses the former, and indirect ones, such as the “citation needed” tag, which addresses the latter (the affordances of the “citation needed” tag will be explored in chapter 4). Conveying malleability also embodies a core tenet of Wikipedia’s epistemology: that it is unfinished (Wikipedia: Wikipedia is a work in progress).

There is a broader aesthetic point to be made here: Wikipedia’s vastness and interconnectedness elicits a sense of wonder. Wikipedia has no start or end: the point of entry is chosen by the reader. As well as not having spatial boundaries, it has no time boundaries either: given the amount of content already present on Wikipedia, and the pace at which it is updated, no single person could ever read it all. This vastness contributes to a feeling of universality and wonder, and to a craving for more that is so familiar to those who have ever plunged into a rabbit hole. When deep in a rabbit hole, it is easy to lose sight of the boundaries of Wikipedia, that frame it as a description of the world, and start to perceive it as a world in itself, similarly to the feeling of being immersed in the virtual world of a videogame. There is a subtle realism to Wikipedia that has to do with its vastness and internal consistency, experienced as truth. It’s hard to keep a critical distance from a cogent system of beliefs, when one is emotionally invested in it, and reaps social rewards from it. Wikipedia’s monumentality, then, becomes in itself a form of validation.

To summarise, in this section I have argued that, by leveraging hyperlinks, the design of Wikipedia’s interface achieves indirectly, and implicitly, addresses concerns around Wikipedia’s sustainability. Namely, hyperlinks connect topical subjects with in-depth information present on Wikipedia, thereby drawing in readers, and implicitly demonstrating the usefulness of the project in its entirety, as a resource to understand the world of the reader. The interconnectedness and autonomy afforded by hyperlinks encourages readers to indulge into Wikipedia’s content, following their interests and curiosity. Wikipedia’s searchability connects the content of the encyclopaedia with the social world of the readers, by affording quick check of facts that can then be spent in social settings, much like traditional encyclopaedias. In doing so, Wikipedia’s interface also validates knowledge by presenting an internally consistent, impressively large body of information, and frames

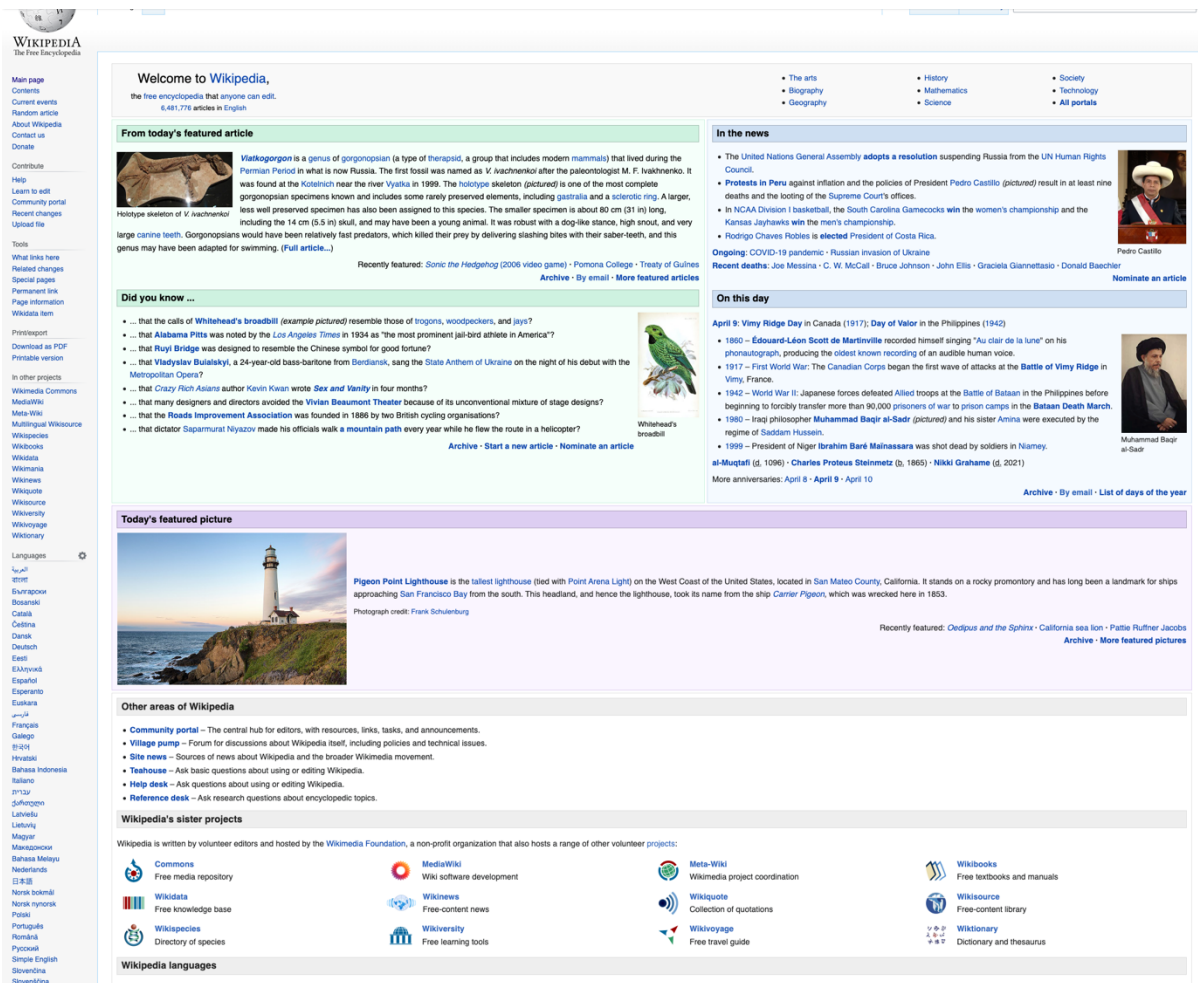


Figure 3.3. Wikipedia's main page.

knowledge as fluid, open to, and indeed needing, constant revision and update.

3.3. Abstraction, situatedness and objectivity

In this section, I will explore a tension between two opposing ways in which users of en.wikipedia.org experience knowledge and knowledge production: as abstract, or as situated. The dominant, normative framing of knowledge brackets materiality and subjectivity; an antagonistic attitude aims at situating knowledge production, both in terms of authorship and sources. I aim to show how design and style choices relate to these framings, by examining the file formats and styles used to represent knowledge, as well as drawing on conversations between Wikipedians and strategy documents published in 2019 by the Wikimedia Foundation, in which strategic priorities for Wikipedia are established.

3.3.1. Abstraction

Several stylistic choices regarding Wikipedia – its content, interface, appearance – result in a framing of knowledge as abstract, by inviting users to experience the content of the encyclopaedia as independent from materiality. Wikipedia's content is presented as distant, as much as possible, from authors and substrates, through design interventions that isolate content from the process of its

production, by hiding authorship, privileging propositional knowledge, and working towards platform-independent content.

I will start with the concealment of authorship, obtained through enforcing a specific writing style, and visually hiding authors from readers. A uniform, unemotional style papers over individual variance, flattening the crowd of authors into an affective minimum common denominator. This effect is obtained by enforcing Wikipedia's house style: the Manual of Style invites Wikipedians to write in a detached, non-judgmental manner. In a section titled Point of view, the manual prescribes what follows:

Quotation should be used, with attribution, to present emotive opinions that cannot be expressed in Wikipedia's own voice, but never to present cultural norms as simply opinionial:

Acceptable: Siskel and Ebert called the film "unforgettable".

Unacceptable: The site is considered "sacred" by the religion's scriptures.

Concise opinions that are not overly emotive can often be reported with attribution instead of direct quotation. Use of quotation marks around simple descriptive terms can imply something doubtful regarding the material being quoted; sarcasm or weasel words such as supposedly or so-called, might be inferred.

Permissible: Siskel and Ebert called the film interesting.

Unnecessary and may imply doubt: Siskel and Ebert called the film "interesting".

Should be quoted: Siskel and Ebert called the film "interesting but heart-wrenching".

(Manual of Style)

What is interesting about this description is the framing: maintaining a neutral point of view is not just framed as a matter of keeping one's opinions to oneself. Neutrality also concerns affect. The result is a rather dry speech, devoid of personality. Emotion shouldn't be brought into writing, thereby concealing an important aspect of the editor's embodiment. In other words, Wikipedia shouldn't sound like it was written by people. I will come back to this point in chapter 4, when I will explain the epistemic role of bracketing the identity of editors. Hiding points of view from a style perspective also means papering over the internal diversity of the editors themselves. Wikipedia's voice doesn't sound like the voice of a collective, but of someone speaking in a monotone.

Visually, authorship is hidden by tucking away the names of authors, in the "Talk" and "History" tabs. This is different from earlier Wikis, where the names of the editors, their comments and conversations, would be mingling with the content, or at the bottom of the page, in ways that are not even consistent: as shown below, examples from MeatballWiki, an early example of the genre – the page about Plato's Republic (Fig. 3.4) has content at the top and comments at the end, the one about Snow Crash (Fig. 3.5) reads like a review or a discussion, but is not signed, and WhatIsArt (Fig. 3.6) is just a conversation.

On Wikipedia, these conversations and names – the engine of article pages – are hidden in the back. In his memoir, Larry Sanger acknowledges this difference, and explains that it was a conscious

choice to leave some parts of wiki culture behind, including the fact that opinions, while they were welcome on traditional wikis, are not welcome on Wikipedia (Sanger, 2005). In other words, points

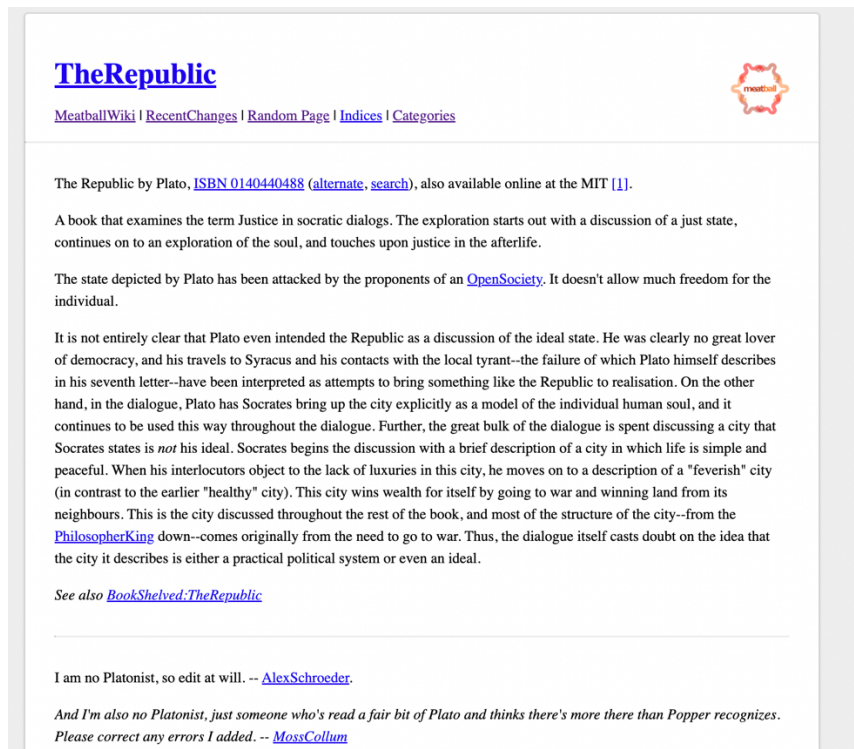


Figure 3.4. MeatballWiki page: TheRepublic.

of view were intentionally removed from the appearance of Wikipedia's pages. I take this to be a form of validation: uniformity expresses a neutral point of view. Wikipedia presents itself as a unified voice, without a point of view.

SnowCrash

[MeatballWiki](#) | [RecentChanges](#) | [Random Page](#) | [Indices](#) | [Categories](#)



From the [JargonFile](#):

[[NealStephenson](#)]'s epic, comic cyberpunk novel is deeply knowing about the hacker psychology and its foibles in a way no other author of fiction has ever even approached. His imagination, his grasp of the relevant technical details, and his ability to communicate the excitement of hacking and its results are astonishing, delightful, and (so far) unsurpassed.

Review by [SunirShah](#): *If you call yourself a webfiend, if you like playing with (vs. "on") online communities, get this book and read it. Do it. Do it now. Or be destined to always be a grainy black and white avatar.*

This book is not only funny, well paced, well written but well thought. Stephenson's obsession with [Wiki:ThreeRingBinders](#), [GatedCommunities](#), and sword fighting twist together to give you a picture of the future as a more characterized today. It's probably going to happen. See you in the [MetaVerse](#).

Review by [DavidChess](#) here: [\[1\]](#) (The first 150 pages were the best, but hey...)

1984, 480pp. [ISBN 0553380958](#) ([alternate](#), [search](#))

More reviews and discussion at [BookShelved:SnowCrash](#)

Spoiler

Just to ruin it for you, the book takes the theory that the human mind is a [TuringMachine?](#) for granted, and then suggests it can be crashed given the right symbolic program; i.e. a [LethalText](#), which in the novel is akin to a worm or virus in the zoology of software. It actually goes much further than this, talking about memetics, the origins of consciousness, the value of individuality, the absurdity of libertarianism, construction of a [HiveMind](#), a rip on Scientology, and a billion other things.

Figure 3.5. MeatballWiki page: SnowCrash.

WhatIsArt

[MeatballWiki](#) | [RecentChanges](#) | [Random Page](#) | [Indices](#) | [Categories](#)



I think art is fundamentally anything that stems from creativity and expression.

"Art is feeling routed through form; and wisdom is emotional knowledge routed through reason." -- [AdinaLevin](#), personal communication to [SunirShah](#), in relation to how sense-making humans differ from animals; e.g. [\[The Symbolic Species\]](#)

Art, I reckon, is the label given to two different strands of the same subject.

The first aspect: Art is an aesthetic. It's about creating beauty. Since beauty is in the eye of the beholder and subject to social norms, art has a fundamentally social character. though obviously the relationship to society can be positive or negative (this idea is where shock art is born).

The second aspect: Art as a powerful communication from one person to another...you know the "it reaches out and touches me" (and not necessarily in a cuddly way) effect. Good art typically makes an effective communication, though the "message" received may vary depending on the viewer: one person sees a painting as an idyllic day of innocent joy on the beach, another person sees the painting as a comment on the vacuous nature of family life etc etc.

Once again, as with anything, there are those who spin off the idea of "art as communication" and create deliberately obscure works. When presented as art, the meaninglessness of the work makes a statement to the viewer... though this would tend to be tinged with one or more of the following: elitism, self referential irony, nihilism etc.

I disagree with the idea of "Art as product of creativity" as it ends up getting sucked down into the romanticisation of lunacy as genius. The elitist idea that only true, credible, individuality can produce art. Since every effect has a cause and nothing that happens is ahistorical then such genius/lunacy/individuality is a fabrication. anyone that existed as far from society as to be considered a true individual would probably not bother with anything recognisable as "art" and if they did would produce something that could not be appreciated by those still trapped inside. -- [OceanHorgan](#)

No art critic would ever claim that individuality is essential to create great art. Most art in the Western world only sought perfection in representation, not individual expression. A masterpiece could be the most perfect rendition of Madonna and Child, although most art served more mundane purposes such as porn or representations of wealth. It's only really in the 20th century that we have seen an emphasis on individuality, and since television, the necessity of the cult of personality in art. -- [SunirShah](#)

Figure 3.6. MeatballWiki page: WhatIsArt.

The second way in which Wikipedia presents knowledge as untethered is by privileging propositional knowledge. Most of Wikipedia's content is written out, as opposed to presented through, for instance, videos or images. While the emphasis on written, propositional knowledge may seem obvious, it really isn't: traditional encyclopaedias have, for a long time, been written on paper, so the medium constrained what could be done. However, when multimedia started becoming a standard, even before the Web became the primary source of information for most people, other formats were explored: Encarta, the CD and DVD based encyclopaedia launched by Microsoft in 1993 (Pallardy, 2019), contained videos and sounds, as well as text. One might argue, of course, that technical requirements must have played a role: producing high-quality videos and images is more costly than writing content out.

It is also true that propositional knowledge makes sense in a setup informed by Western knowledge (Van der Velden, 2013), and fits with the traditional, written form encyclopaedias were typically found in at the time. I agree that cultural and historical factors play a role here, but I would argue that, in addition, propositional knowledge fits with Wikipedia's specific understanding of epistemology, and of its own role in knowledge production.

Propositional knowledge is important to Wikipedia also in an ideological sense. It is important enough to be part of the rationale for Wikipedia's logo (Fig. 3.7). The logo represents a globe made up of puzzle pieces, each of which presents a character in a different alphabet. As outlined in the official visual identity guidelines:

The Wikimedia Foundation marks represent much of what our projects and our movement stands for: bold ideas, global collaboration, typography and text, integrity and quality.

(Wikipedia's visual identity guidelines, Wikimedia Foundation)



Figure 3.7. Wikipedia's current logo.

Typography and text are not just means of expression of content, they are framed as values in and of themselves, listed alongside universality (“global collaboration”) and epistemic values (“integrity and quality”).

There are intertwined technical affordance and ideological reasons why propositional knowledge is so fundamental to Wikipedia, and, in my view, they have to do with Wikipedia’s values around collaboration and universality. Wikipedia is founded on an ethos of radical collaboration – anyone can edit anyone else’s work. We must keep in mind, as well, that Wikipedians collaborate in a decentralised manner. The ideal format needs to allow for small changes to occur independently from one another. Text works well: it is composed of units (paragraphs, sentences, words, etc.) that can be easily tweaked in small increments. Doing the same with an image, or a video would simply be impossible, if one wants to preserve some level of intelligibility. Text also has a relatively low (but not the lowest) barrier of entry: those who are envisaged as potential editors are people who, already, had the ability to read and write – but perhaps not draw or edit video. Granted, this is culturally and historically situated: there have been times, in the West, when drawing was considered an essential skill, part and parcel of knowledge production (see, e.g. Reinhart, 2022; Ambrosio, 2014; multiple examples in Daston & Galison, 2010).

Secondly, text encodes ideals of universality – not because it can be read or written by everyone (it can’t), but because it is transportable. Text can travel: before the Internet, it could be copied, or dictated across distance. Even with contemporary technology, text is still more transportable than other formats, as it is lighter, in terms of data, than images or videos. Text can also be translated: even accounting for discrepancies in terms of coverage (languages having words for things other languages don’t have words for, differences in meaning and reference), the bulk of what is and can be said can be translated.

Text is also easily replicable. One of the principles on which Wikipedia is built is the ability to fork: that is, to duplicate Wikipedia data (with the exception of user data) and start a copy of Wikipedia elsewhere, known as a “mirror site”. Historically, forkability has been important as an insurance against hardware issues, and indeed, Richard Stallman, in his project for a free encyclopaedia, recommends to allow mirror sites:

Permit mirror sites.

When information is available on the Web only at one site, its availability is vulnerable. A local problem—a computer crash, an earthquake or flood, a budget cut, a change in policy of the school administration—could cut off access for everyone forever. To guard against loss of the encyclopedia’s material, we should make sure that every piece of the encyclopedia is available from many sites on the Internet, and that new copies can be put up if some disappear.

(Free Software Foundation, 2024)

Stallman wasn’t writing about Wikipedia, but his ideas for a free encyclopaedia were, as

described in chapter 1, extremely influential on Wikipedia, hence an ideological lineage can be easily traced here. Should Wikipedia's servers shut down, or be attacked, mirror sites would be able to compensate. Forking comes from programming practice, and concerns the ability to copy existing code and use it or build upon it. Forking only makes sense within open software culture: proprietary software can't be copied – hence it can't be forked (Tkacz, 2011).

Over time, forkability, from being a security measure, has turned into something closer to a way of legitimating power: if disagreement over editorial choices becomes deep enough, anyone can, in theory, download Wikipedia's data and software and start a rival encyclopaedia elsewhere. This is what happened when Spanish Wikipedia forked (Reagle, 2010). Forking is, however, much more difficult in practice than in theory, for reasons related to skill and computing power (Konieczny, 2009). Hence, forking constitutes a largely theoretical form of dissent, that can nonetheless be held against those who disagree, in the form of "if you don't like it, you can always fork" (Tkacz, 2011; Tkacz, 2015).

Forkability, then, is fundamental to Wikipedia. It also influences its writing style – text needs to be adaptable. For instance, the content of articles cannot be self-referential, because mentioning Wikipedia would make the text unsuitable for re-use within a different project. Wikipedia's Manual of Style upholds the principles described above:

Mentioning that the article is being read on Wikipedia, or referring to Wikipedia policy or technicalities of using Wikipedia, should be avoided where possible. This type of self-reference limits the use of Wikipedia as a (Konieczny, 2009) encyclopedia suitable for forking, as permitted by our license. The goal of Wikipedia is to create an encyclopedia, not merely to perpetuate itself, so the articles produced should be useful, even outside the context of the project used to create them.

(Wikipedia: Manual of Style/Self-references to avoid)

This prescription points beyond self-reference, and towards a more general principle of transportability and independence from its source. For content to make sense on a website other than Wikipedia, it has to be written in a way that does not explicitly depend on the platform where it's published. There is a spectre of free software ideology here as well: the product needs to be such that it can be replicated and reused in different contexts; hence it should be as untethered as possible from its context of production. Requiring transportability and independence from source also constrains the choice of format. Text is a good candidate in both cases: as mentioned above, it is light on data, hence easily transportable, and, intelligible in a range of material contexts.

Pressed by the evolution of the hardware market, and calls to include non-Western knowledge, Wikipedia has also, at various points, and with very limited success, experimented with audio. From a technical point of view, the spoken word still affords some of the uses described above: it can be translated, and it travels relatively well (but not as well as written knowledge). On the other hand, spoken language doesn't afford collaboration in the same way as written language: it would be very strange to have articles in audio format, where each word was spoken by a different editor; audio editing is also not a widely held skill, and is not easily supported by Web pages, at least in their current form. There are ways of collaborating on audio, as, for instance, happens on TikTok, where

users can take a sound made by another user and build on it, but they are hardly conducive to creating an encyclopaedia in the traditional sense.

Using audio in relation to oral knowledge, as a way to include modes of knowing that are not Western – from places where oral transmission is privileged, for instance – is riddled with epistemic complications. One of the very few attempts at including oral knowledge was the Oral citations experiment, which took place in 2014/2015, and comprised of three small articles containing information derived from interviews (Wikipedia: Oral citations experiment). The experiment was short lived and only involved articles about topics that relate to indigenous people: Otjiene (a village in Namibia), the Herero people, and Kuaima Riruako, a Herero politician. The information included in these pages, obtained from interviews, is about local lore and history. Factual information is entirely drawn from written sources.

This short-lived experiment highlighted a different aspect of the importance of propositional knowledge, and that is: provenance. Oral knowledge is not just transmitted orally: it's also validated in a different way from written knowledge. Oral sources in the Oral citation experiment support essentially unverifiable information, framed as narratives. Accepting oral knowledge as a support for factual content would necessitate amending the verifiability policy, that only accepts published sources. The framing of knowledge as written, propositional, and mediated by a specific structure of validation, has been inherited by Wikipedia in the form of academic conventions. There might also be a case to be made for the framing itself to have a role in this: for those Wikipedians who have a strong attachment to Western conventions, supporting verifiability might also be a way to oppose the extension of epistemic parameters, so that they might include knowledge validated in other ways.

When it comes to the issue of transportability, there is a further complication. Early Wikipedians upholding forkability could assume that Wikipedia's content could be re-used, and trusted, elsewhere, without being tethered to sources, because of an assumption of validation: given the majority Western setup of the Web, and encyclopaedias, one could assume that information had been collected from sources considered reliable, and not, for instance, just made up. In other words, Wikipedia's mirror sites could count on an implicit, assumed provenance, and trust it, in part, because of an implicit trust in the default Western system of validation. Diversifying the ways in which knowledge is validated would entail framing information differently when it's transported: in the case of the Oral citations experiment articles, for instance, some kind of sign that information was derived from interviews and local lore would be needed, thereby weighing down free-floating information with provenance, which is contrary to the spirit of forkability.

Finally, abstraction is enacted as a technical value in the form of platform-independence. Platform-independence has become important because of the role of mediation in the construction of the idea of abstraction, as I explain in what follows.

Propositional knowledge is, in part, experienced as abstract because, in the context of the Web, the written word is transparent. A technology is transparent when blends with the experience of its use, to the point that users stop thinking about the artefact itself, and only focus on what they are doing with the artefact (see Verbeek, 2005; Ihde, 1990; Heidegger, 1977 for more on transparency). Transparency fails – and the object becomes noticeable again – when something breaks (Verbeek, 2005). For instance, as I type this, I am not thinking about the keyboard of my laptop – I am focusing on the meaning of the words I am typing. Should, say, the “c” key stop functioning, I would be

painfully aware of the existence of the keyboard, and of the materiality of the hardware I am using. Transparency is predicated on the absence of friction.

The written word on the Web enjoys a level of transparency because it can be easily manipulated with available technologies, and through widespread skills. Websites are primarily text-based – even more so when Wikipedia was created – and sit within an ecosystem of technologies created to handle the written word. Since most people don't think about the devices used to write or access written knowledge (keyboards, screens), the pipeline of written propositional knowledge enjoys a level of transparency. When a new format is introduced, small or big adaptations need to be concocted, concessions made, workarounds found – friction is back, and with it, the awareness of materiality. The seamlessness of the written Web creates the collective illusion of the absence of materiality – of abstraction. Abstract is what coincides with the available.

Innovation in hardware and modes of consuming content have, by introducing a wider range of mediations, forced Wikipedia to consider diversifying its output. The WMF is now pondering how to create a product that can be consumed across platforms:

We might start doing experiments of what does it mean to do that – to send pieces of data, like we're talking about like a knowledge store.

Right, like knowledge as a service, what does that mean when it's not a Web page. Facts are probably too small, but a webpage is too big like what would something like your smart TV or Alexa or whatever need to get in order to make like the connections and responses.

(Moriel Schottlender, *Principal Software Engineer, Wikimedia Foundation*)

The white paper titled *Experience* is also clear on the necessity to keep up with multiple ways of consuming knowledge, and recognises that the format of Wikipedia content is not up to scratch.

A significant challenge in making all the world's knowledge accessible to all the world's people will be to ensure that it is optimized and future-proofed for a rapidly-evolving digital consumption environment. The term Form Factor generally refers to the various entry points, devices, channels and formats that define a digital product offering. In the context of Wikimedia, form factor will include (a) the variety of devices that Wikimedia content can show up on now and in the future; (b) the size, and flexibility of the content itself.

(*MediaWiki: Experience [white paper]*, p. 32)

A major worry concerns social media:

social media channels have (re)defined the expectations and habits of users all over the world. Users now expect relevant information to find them as a result of their preferences, feed settings and serendipitous browsing. Facebook

[3] and WhatsApp have become primary entry points for new users accessing the Internet, and are, for many emerging communities, simply conflated with "the

Internet”. For these communities, the page-based mental model [4] of the Internet will effectively never have existed.

(MediaWiki: Experience [white paper], p. 61)

A potential solution to the need of making Wikipedia’s content work across platforms is to chunk it up into units (“facts”), so that they can be recombined and transferred through multiple mediums.

All of the scenarios above (adapting Wikipedia content to multiple devices) are problematic for us right now, mostly because of one thing - our core, fundamental element is the article, not the fact.

As mentioned in the Discovery document, major Internet players like Google and Facebook are already grabbing our text content, chopping it up, and presenting it in factoid-sized chunks. We currently don’t have influence over this process and the resulting user experiences, largely because we don’t have any facility that we can point to and say, “do it this way, it’s better and we’ve already done the work for you.”

There are several ways we might achieve this “atomizing” of articles:

- Automatically break up the entire article into elemental parts (sentences/passages perhaps)
- Take the top 5 most important elements/ sections from each article and atomize that (although it’s unclear at the moment how we would identify those top 5 elements).
- Have the community decide which elements should be atomized for each article (essentially a new editor function/workflow).

(MediaWiki: Experience [white paper], p. 36).

In the same white paper, atomised content is deemed more suitable for use in, for instance, audio mediums, such as Alexa.

Another way to achieve platform independence is to uncouple content and format – that is, to make the content more and more abstract. Part of the rationale for abstraction is ethical. Connectivity, and devices used to access the Internet vary geographically. Since different formats work on best on different devices and available technologies, and different technologies correspond to different levels of connectivity, freeing information from format might allow for wider availability of Wikipedia’s content, reaching underserved populations.

In Wikimedia Foundation documents, expanding the geographical reach of Wikipedia through transcendence is sometimes labelled “ubiquity”.

we consider this issue of ubiquity at the intersection of Wikipedia and its consumers, structured content would relieve us from the requirement to anticipate, monitor or otherwise be directly aware of how all populations in all

emerging economies are developing their own unique relationships to the Internet.

(MediaWiki: Tools [white paper], p. 92)

Abstraction is explicitly held up as an ideal, in the form of abstract Wikipedia, a project that aims to untether Wikipedia content from language:

The goal of Abstract Wikipedia is to let more people share more knowledge in more languages. Abstract Wikipedia is a conceptual extension of Wikidata.^[1] In Abstract Wikipedia, people can create and maintain Wikipedia articles in a language-independent way. A particular language Wikipedia can translate this language-independent article into its language. Code does the translation.

(Abstract Wikipedia [project page])

Abstract Wikipedia rests on a project called Wikifunctions, a repository that provides code for anyone to use, including code to translate Abstract Wikipedia data into any language that is already associated with a local edition of Wikipedia. Abstract Wikipedia is an ideal to tend to – and the practical steps being taken to attain it are, predominantly, to build and populate Wikifunctions. If Abstract Wikipedia were to become Wikipedia’s predominant way of operating, then the merging of coding and writing Wikipedia would be, essentially, perfect: Wikipedia would consist of a database that can be then turned into content in multiple ways, and propositional knowledge wouldn’t necessarily constitute the dominant paradigm anymore.

Another way of achieving platform-independence is to move away from writing. As sketched above, however, moving away from writing means thinking about the system of validation as well: if a given fact is, say, shared via an audio medium, how can provenance be expressed? Of course one can mark information with an interruption, but when encyclopaedic content of the kind published on Wikipedia, which is ideally entirely referenced, constant interruptions for references can make the listening experience rather jarring.

The problem of platform-independence highlights how important aesthetics is, to Wikipedia: it uncovers how Wikipedia is not just a collection of facts (a fantasy that some parts of the community have been cultivating since the beginning), but also a means of entertainment, as encoded by the shop-window aesthetic of the main page, and the cultural practices around the use of Wikipedia, as described in section 3.2. As highlighted in section 3.2, the element of discovery is an essential constituent of Wikipedia: the experience of reading Wikipedia is as important as the outcome of reading Wikipedia – i.e. learning something.

As of March 2023, the discussion around audio formats is still live. However, Wikipedia now has its own sound (an audio logo, in marketing lingo), the “sound of all human knowledge”. It was selected through a contest between Wikipedians, and won by a Thaddeus Osborne (Wikimedia Foundation, 2023). Primarily, this sound is not about provenance, at least not directly (one might say: it is a stamp of quality because the sound is a reference to the fact that the source is Wikipedia, so the validation of whatever fact comes before or after the sound is rooted in Wikipedia’s knowledge production process). The Wikipedia sound exists for branding purposes: it’s a way for Wikipedia content to be recognised as such. The need for a mark of recognition is rooted in worries around

Wikipedia content being repurposed and used by other companies without marking provenance.

Privileging abstraction is consistent with Wikipedia's ethos and design: it tallies with its core policies, permits to enact core values, matches with the cultural landscape in which Wikipedia is inscribed. Abstraction is expressed through the use of symbolically light means of expression, such as text and language-independent data, and by aiming to detach the encyclopaedia's content from the materiality of its substrate, as is the case in the ambition for platform-independence. Concessions to the abstract ideal, however, appear to be necessary, beyond the trivial observation that some kind of material substrate is necessary for the propagation of information. The illusion of abstraction that was maintained, to an extent, thanks to the adherence between Wikipedia's content and the site of its production, has been dissipated by the evolution of the hardware market. The disappearance of substrate transparency pushed the desire for immateriality further, in the form of platform-independence. Worries around relevance, which inform, to an extent, the abstract ideal also, on the other hand, inform the opposite strategy, which I will describe in section 3.4: situating information in specific ways, that align with new forms of embodiment.

3.3.2. Situatedness

In this section, I am going to describe a backlash against Wikipedia's tendency to abstraction, in the form of a desire for situatedness, coming in part from outside, in part as a result of internal pressures related to worries around relevance.

I noted in the previous section that Wikipedia's content is presented in a uniform, minimalist style. The problem with this style, put bluntly, is that it's boring; or as euphemistically put in a strategy document:

If you're the kind of person who loves deep dives into complex topics, and you don't mind spending time with text that can be challenging, Wikipedia's current format totally works for you.

(MediaWiki: Experience [white paper], p.20)

As explained above, one of Wikipedia's main worries, at the moment, is to appeal to a continuously growing audience, whose taste is changing. In the white paper quoted above, multiple references are made to the fact that new generations of Internet users are accustomed to enticing, multimedia content. In comparison, longform articles might make Wikipedia look obsolete – and worse: dull.

The worry around the length of articles was already present back in 2003. The following is again from early Wikipedia mailing lists conversations:

If we want Wikipedia to be open to everyone, and easy to edit, I think we should seriously consider aiming for shorter articles everywhere. A reader who wants to read 30K of information about a subject would still be able to; they'd have to read three articles instead of one, maybe, but it would only involve two clicks of the mouse...

(Oliver Pereira)

The length of the articles was also related to their scope:

A subject which cannot be covered in such detail should not have its own article. It should not be merged into another text if it can realistically grow to that size when relying on verifiable information. I think you understood it this way, but I'm not sure everyone else did. The length itself is debatable, but I think most of us agree that we don't want one-paragraph texts about every single fictional character out there.

(Erik Moeller)

This quote is followed by a lengthy list of arguments for and against long articles. Both quotes are from a long conversation on this topic, stretched over 46 emails.

Multimedia is seen, as well as a way to re-package Wikipedia content so that it's suitable to be consumed on multiple devices, as well as a way to make content more aesthetically appealing. In the white paper *Experience*, the authors imagine ways in which Wikipedians can collaborate on multimedia content:

Our strength is the Wiki way - allowing multiple perspectives and contributors. Fortunately, we don't have to guess or theorize about what this might look like. There's a perfect real-world example in the form of *Winter on Fire*, [10] an Oscar nominated, Emmy nominated, feature length, Netflix original documentary that chronicles the deadly anti-government protests in Kiev, Ukraine that took place in 2013.

Winter on Fire had 28 credited cinematographers, using video footage captured by ordinary people who were on the ground during the conflict. In many ways it was a crowd-sourced film. [10]

(Wikimedia Foundation, 2019, p.16)

This quote is interesting because of how it frames collaboration and objectivity. There are many ways of collaborating on a video: indeed, most video content out there is produced by teams of people. Typically however, the team is composed of people with different roles – cameramen, photography, sound, etc. – who coordinate among themselves, often through specific roles, such as that of the director, on a shared vision.

The way collaboration is envisaged in this white paper is different, and modelled on the way Wikipedians already collaborate on text. Rather than each person having a role, each person makes their own little segment of video, doing everything necessary; the video is then, rather than an organic representation of something, a composite. To an extent, traditional videos are composites as well (cameras looking at different parts of a scene, and then perspectives are edited together through montage); their narratives, however, aren't.

Interestingly though, the result of a video composed of multiple little videos, as the white paper envisages, while being modelled on the process of Wikipedia's textual production, would create the opposite effect to the uniform, monotone drone found on the page. Instead, it would root each contribution into a point of view, arching towards a form of objectivity that is the opposite of

neutrality – rather, it would be literally the sum of different points of view on the same subject.

A minor point on situatedness can be made about the Wikipedia sound as well. While, as pointed out earlier, using sound is a step towards untethering information from its source, at the same time, the specific sound chosen to represent Wikipedia carries specific referents.

This is the inspiration of the sound, according to the Wikimedia Foundation website:

It combined the sounds of turning the pages of a book with keyboard clicks and a synthesizer chime, representing the huge body of knowledge that people can access digitally through Wikipedia and other Wikimedia projects.

(Wikimedia Foundation, 2023)

The sound is the sound of things that people use or have used to produce knowledge, and references a specific way of producing knowledge. Once again, written sources – book, whatever the keyboard is accessing or producing – are the main shorthand for knowledge production, despite the challenge of translating them into sounds. If oral knowledge had been of any importance, human voices could have featured. In any case, the sound is the sound of specific objects, that indicate a specific, situated way of producing and accessing knowledge.

Situating knowledge – aesthetically, by expressing it in context-specific forms, and epistemically, by rooting it in specific points of view – seems to be a plausible option for Wikipedia’s future development. Situatedness, however, sits uneasily with Wikipedia’s culture, creating frictions with the implementation of a neutral point of view, and explicitly appealing to future readers, who are framed as different from the existing community – hence, by definition, an out-group. I will return to this point in later in the chapter.

3.3.3. Abstraction, situatedness & justice

The tension between abstraction and situatedness, within Wikipedia, is important in a few different ways. I have outlined above how it is embedded in technical priorities, and how it relates to the technological and cultural landscape within and around Wikipedia. Framing knowledge as situated has, historically, been a staple of feminist epistemology (Haraway, 1988; Harding, 1992; Hill Collins, 2000). Justice-oriented epistemologies have long aimed for situatedness, one way or another. In Wikipedia’s case, situatedness has been leveraged in order to pursue some forms of epistemic equity, in ways that tie with the feminist framing described above. Interestingly though, so has abstraction. In this section, I will describe two proposals, both aimed at improving Wikipedia’s fairness, but in opposite ways: Heather Ford (2022) calls for a stronger tethering of information to its source, while Maja van der Velden (2013) suggest doing the exact opposite.

At the end of her book *Writing the Revolution* (2022), Heather Ford draws conclusions from her description of the way in which the article “The Egyptian revolution of 2011” took shape. She examines the interplay between conversations between editors, Wikipedia culture, infrastructure, and the wider technological and social landscape in which Wikipedia is located. At the end of the book, she concludes that history, on Wikipedia, is written by a specific subset of people, and that this bias is influenced by the way knowledge is produced:

The resulting representations are not an indication of everyone's memory or meaning of what happened. Particular types of actors are more suited to the rapid retelling of events in real time and according to the logics of data. The form and temporality of the representations that dominate our knowledge of what have implications for who has the power to influence them.

(Ford, 2022, p. 135)

She identifies a few issues with the process of making knowledge on Wikipedia. One of them has to do with obscuring the provenance of facts:

When facts are translated into semantic data and moved from their source, they shed the traces of their origins—including their companions (sources) as well as traces of the decision-making processes by which they were made.

(Ford, 2022, p. 136)

Her solution, then, is to highlight provenance, by situating data through design interventions:

Situating data requires embracing design principles that foreground the ways in which facts were produced and by whom they were produced and highlighting the volatility of information affected by historic events. Knowledge graphs that display a single answer to users' queries, often without the source being visible, can highlight the origin of facts and their paths to the knowledge graph in ways that surface the knowing subject. This includes finding ways of attributing both the source and the rules or assumptions that are made in determining the prioritized, single answer to users' queries. This may seem unwieldy for a form that aspires to simplicity, but there will certainly be creative ways to design such reflexivity into automated systems.

(Ford, 2022, p. 138)

Ford's reasoning here seems to be that awareness of sources of information and the process of construction of knowledge is, in and of itself, a form of justice, as, rather than representing Wikipedia's content as purely neutral and objective, it highlights its contingent, constructed nature.

Maja van der Velden, in *Decentering Design: Wikipedia and Indigenous Knowledge* (2013), argues for the opposite solution: further untethering. Her starting point is the representation of knowledge: she asks whether indigenous knowledge can be represented fairly on Wikipedia. Her answer is that Wikipedia is designed to reflect a specific organisation of knowledge and, as such, struggles to fit non-Western systems within its infrastructure:

Wikipedia's design does not allow for Indigenous communities to use Indigenous concepts and structures to tell a story and to present and organize knowledge. At the same time, articles about Indigenous knowledge, artifacts, and peoples are categorized in a manner that fragments Indigenous knowledge and makes it invisible as a body of regular, rational, and systematic knowledge.

(Van der Velden, 2013, p. 311)

Van der Velden's solution is to make Wikipedia's structure disappear, thereby tending to extreme abstraction:

It is [...] possible to imagine Wikipedia as an ontologically flat database, a database with millions of articles but without a central design to organize the items and to make them accessible. Different interfaces would enable different ways to access to the same body of articles. Some users would like to access Wikipedia through its current search or browse options, whereas other users would prefer a design closer to their understanding of the world. Second, different pillars and protocols would enable Indigenous knowledge communities to use different media and oral citations to record their knowledge.

(*Van der Velden, 2013, p. 314*)

According to Maria van der Velden, abstracting from categories, and transforming Wikipedia into something like raw material, can have positive consequences on knowledge inclusivity. Her proposal is very close to Abstract Wikipedia, described above; although the code itself would still be written by a specific group of people, in such a way that might limit the options for display. While the content would be abstract, the design of the system would still be situated, and it seems hard to imagine how it could be any other way.

The comparison between these two proposals is interesting for three reasons: first, in terms of what it means to design for epistemic justice; secondly, because they reflect different conceptions of what justice is; third, how they reckon with the materiality of Wikipedia as an artefact. The first observation that can be made is: there isn't only one way of turning epistemic justice into design. This observation may seem obvious, and perhaps inconsequential, but it will become important when I will talk about the process of value embedding, in chapter 5. In this case, two opposing movements – towards and away from connecting a fact or data to its source – can both be, reasonably, argued to contribute to epistemic justice. In Ford's case, provenance improves our understanding of the fact that knowledge isn't universal, thereby highlighting the role of power in creating knowledge, and instilling a healthy diffidence in the reader, who will hopefully be inclined to consume Wikipedia content critically. In Van der Velden's case, knowledge that is presented according to modes that are native to the dominant group is made malleable enough to fit alternative ways of understanding the world, thereby empowering the reader to use their own categories in order to make sense of the data they are presented with. Both proposals are internally consistent; the design suggestion depends on what exact conception of justice each author is working with.

Which brings me to the second point: justice is not a sufficiently specified concept, when it comes to designing an artefact. It is necessary to ask, for instance, what kind of justice we are looking for, and for whom. Ford is mostly inspired by feminist epistemologies, which moves from a critique of objectivity as the point of view of the dominant group, concealed under the pretence of universality. Hence, a partial solution is identified as unconcealment, and consequent questioning of knowledge in and of itself. Van der Velden, instead, operates from a Postcolonial Computing perspective, frames design itself as a site of power, under the assumption that design is always from somewhere (Van der Velden, 2011) and seeks to decentre designers in order to achieve equity (Van der Velden, 2013). From this perspective, the site of power is infrastructure, and as such, its design

needs to be delegated as much as possible: hence the flat database, that users can configure as they see fit. Interestingly, both Ford and Van der Velden use standpoint epistemology in their work, and start from a similar premise: that knowledge is situated. The final, technical result however, diverges.

Finally, one word on materiality. Ford and Van der Velden are analysing the same object – en.wikipedia.org – but focus on different parts. Ford focuses on the process of knowledge-production, and the links between Wikipedia and the rest of the Web, specifically those devices that allow for facts to travel (e.g. Google, hyperlinks). Van der Velden, on the other hand, is mostly concerned with categories, internal to Wikipedia: that is, with those sections of the website which allow for knowledge to be boxed up and linked. Of course, one could argue that the different focus is justified by their respective interests and theoretical backgrounds, and that would be correct. It is also, though, important to note that once that choice of focus is made, the range of possible interventions is restricted, and as such, the embodiment of values such as situatedness and abstraction. Once homed in on knowledge-production, Ford has little choice but to advocate for situatedness: given the specific portion of Wikipedia she is looking at, abstraction would not support justice in any meaning of the concept. The system of links and sources affords situatedness as a possible solution to the problem of justice: it is possible to make provenance more obvious, while it's not, for instance, possible to force editors to use different sources. Van der Velden has a bit more latitude, as she explicitly advocates for doing away with categories all together: but again, the way Wikipedia is built – on propositional, chunkable knowledge – is such that a flat database might work. Paradoxically, had Wikipedia contained a more variable range of formats, such as audio and video, which are often advocated as ways of decolonising Wikipedia, the idea of a flat database would sound less doable.

Highlighting provenance and situating the production of knowledge seems to me to work against the principle of forkability as well: Wikipedia's content can hardly be replicable and modifiable while maintaining ties to the context of its production. On the contrary, it has to be created in a way that actively discourages foregrounding the process of knowledge-production, for ideological reasons that are embedded in Wikipedia's functioning. Here, I am not advocating for any of the factors above to come first: rather, I suggest that an interplay between different conceptualisations of justice and materiality ultimately circumscribes design interventions. I will return to this point in chapter 5.

The relationship between information and the substrate hosts, on Wikipedia, a space for the deployment of conflicting ideals of knowledge. Abstraction from materiality, when inscribed within Wikipedia's understanding of objectivity as neutrality aligns with a framing of knowledge as rational, disembodied, and independent from its source; when evoked in the context of postcolonial epistemic practices, it can be leveraged to support a commitment to a variety of epistemologies. Situating knowledge by highlighting its provenance, on the other hand, undermines its universality, thereby positioning itself squarely against the notion of objectivity as neutrality. Analysing the relations between epistemic commitments and possible implementations of Wikipedia's content highlights the complex entanglements between ideology and technics. In the following section, I will explore how design choices, discussed in relation to imagined future users, contributed to define who belonged to Wikipedia's community.

3.4. Belonging

Design is often a marker of belonging to a given community. Through the construct of taste, owning and using certain objects often projects one's belonging to a certain group of people, characterised by a certain sensibility, status, view of the world (Bourdieu, 2002; Parsons, 2016). Conversely, artefacts are designed with a certain audience in mind, bringing in issues around belonging, through the visual appeal of objects, and shared values.

Early conversations between the group of people who created Wikipedia display a preoccupation with aesthetics as it relates to universality, in the form of the – unspoken, gestured towards – question: who should Wikipedia appeal to? Given Wikipedia's ambitions to universality, the obvious answer is: everyone. The question, for the founders, was: how do we make Wikipedia as widely appealing as possible? The answer wasn't obvious, and reading those early exchanges, it becomes clear that participants in those conversations weren't sure who "everyone" was.

During those conversations, Wikipedia's early designers were creating a shared understanding of, and debating, who the user of the product was. The definition of the user occurred in relation to themselves: the conversations can be read as discussions of similarities and differences between the prospective user and the designers themselves. In design theory, the fictional person or people the early designers of Wikipedia were discussing is called the "imagined user", i.e. the person or people an object is designed for, as they are represented by the designers (Ivory and Alderman, 2009; Ivory, 2013). The imagined user is a useful construct because it provides a focus to examine the relationship between design choices about Wikipedia's aesthetic and issues of belonging and community. In the following, I will outline two different kinds of imagined users – newbies and skilled editors – as they are constructed through conversations around the design of Wikipedia's skins. The way these two classes of users are conceptualised, I argue, has influenced the development of Wikipedia's interface.

An interesting negotiation over design choices occurred, in the early days of Wikipedia, around the customisability of Wikipedia's interface. Wikipedia is a skinnable website, meaning that readers can choose what skin (external appearance) they want to use when reading Wikipedia. Skins have been developed over time, and however much freedom is granted to the user, they still require some kind of design choices: first of all, the skins need to be designed, so decisions are made at that point; then, it needs to be decided what the default option is. These choices are made on the basis of various factors, including the imagined user. Skins can be created and downloaded by users as well. In this way, Wikipedia looks whatever people want it to look. Even for those who are not able or interested in using skins, Wikipedia's HTML is not very prescriptive: for instance, it doesn't set a font. It only sets "sans serif" (the simplest, least frivolous choice), meaning that the user's browser will select its default sans serif font, which can vary depending on the device's operating system.

Early conversations around skins make an explicit distinction between an in-group, of established members of the Wikipedia community (which includes the participants in the conversation) and an out-group of potential editors. The following is about the first kind of imagined user, the skilled expert, based on the participants in the conversation. The similarity between this kind of imagined user and the participants is marked by references to what is obvious and common sense – that is, that which aligns with the beliefs and habits of those who are speaking.

An important value expressed over and over again is that of freedom of choice, expressed through

technical features such as default and customisation:

I think that saying it is "a new default skin" is not the way to go. Rather, first implement it as an alternative skin, and only then start seeing whether people want it as the default skin.

(Andre Engels)

What precisely is terrible about it [the existing skin]? IMO... - It leaves the user in control of font family and size, which is very good. - It is white-on-black, which is too contrasty for optimal reading for me - bad

(Richard Grevers)

Leaving it up to some imagined users, those who were technically skilled enough to, for instance, change the default font in their own browser, to make choices about the appearance of the website is considered a good thing:

In IE, Tools, Internet options, click the "fonts" button on the first panel: "The fonts you select here are displayed on Web pages and documents which do not have a specified text font" - it's generic as can be. And I have 13pt Trebuchet selected in every browser

(Richard Grevers)

This contrasts with the way appearance is talked about when newbies are the imagined users, when the level of expected skill – and the extent to which the user is expected to simply use what they are given – is different. Newcomers are imagined as low-skilled users who use the default:

Our default skin is terrible. We have to implement something which is easily usable by the masses who read Wikipedia. Let's work on a good clean, cross-platform and -browser design which does not turn newcomers away.

(Tarquin)

we *badly* need a more noob-friendly skin. (For example: I have been trying to get Wikipedia bookmarked in all libraries across my area, but the way it looks at the moment, most people just find it too confusing -- people have to remember: users start reading at the top left and keep going. At the moment, they think our site is called "Main Page <http://www.wikipedia.org/wiki/Main_Page> | Recent changes <<http://www.wikipedia.org/wiki/Special:Recentchanges>> | Edit this page <http://www.wikipedia.org/w/wiki.phtml?title=Main_Page&action=edit> | Page history <http://www.wikipedia.org/w/wiki.phtml?title=Main_Page&action=history>")

(Erik Moeller)

I don't think that what registered users like is as important as what is sound design and clear for unregistered users. Check my earlier post on "Lusers' eyes glaze over when faced with Wikipedia pages".

(Erik Moeller)

In these messages, the in-group (those who are working on developing Wikipedia, and, to an extent, registered users) is clearly demarcated, as opposed to an out-group of imagined users – those who may or may not want to join Wikipedia as editors. Non-Wikipedians are framed as people who don't have technical expertise: "noobs" and "lusers" both meant, in online slang at the time, people who didn't have expertise in technology, or anyone who was not a geek. All the labels applied to imagined users, in these conversations, are rather condescending: "noob" has a similar meaning to "newbie" – initially, in early tech culture, it was applied to those who were new to coding (and often asked "stupid" questions in forums); "luser" is self-explanatory, given its assonance to "loser"; so is "the masses". One kind of imagined user, then, was a relatively low-skilled person who could nonetheless be useful to the project. Concern around losing editors was, clearly, already present at the very beginning (this thread is from 2003, two years after Wikipedia's launch).

Customisation separates people in two groups in another sense as well: being able to customise technical artefacts has historically been a big part of FOSS culture – indeed, the idea of taking and repurposing software is fundamental to the FOSS ethos. Customisation also implies a trade-off: on one hand, a coded artefact that can be easily customised offers a wider range of options to the user, while a ready-made, boxed one doesn't; on the other, customisation does require technical skill, while ready-made software doesn't. So customisation, in and of itself, operates a sorting operation between groups. Sorting users in this manner happens frequently on Wikipedia; I will provide further details in section 4.3.2). Customisation also embeds values of freedom and self-determination that are dear to the FOSS movement.

Another interesting case study, as far as probing the politics of belonging and design choices goes, is the history of the Wikipedia logo. There are two things going on here: the history of the logo itself, and the way Wikipedia represents it.

While reading early conversations between Wikipedia founders, I came across the following quote:

PS: I'd recommend you change the American flag logo. Extremely ethno-centric et. al. I think a globe logo would be much more fitting, if you want to keep with that metaphor. Or perhaps a book.

(Scott Moonen)

This immediately attracted my attention, because I had, thus far, never heard of the American flag being a Wikipedia logo. The page "Wikipedia logos" (Wikipedia: Wikipedia logos) features, as first logo, the one designed by Bjørn Smestad, a sphere with text wrapped around it; the text is a Lewis Carroll quote. There is no mention of an American flag. A clue to the existence of other logos, that are not shown on the page, is in the text, which describes it as Wikipedia's first "official" or "true" logo, without expanding on any "unofficial" logos. Eventually, I found the page "Logo/History" (MetaWiki: Logo/History), shown in Fig. 4.6, which does show an American flag as the first

Wikipedia [edit]

Progression of logos [edit]

Image	Title	Years in use	Description
	Test logo	2001	In the early days of Wikipedia, founder Jimbo Wales uploaded an American flag as an example logo. Needless to say, it was never intended to be the permanent logo. See the following discussion at OldWikiPediaLogo .
	Nupedia logo	2001	The Nupedia logo was originally sent to a Nupedia logo competition by Bjørn Smestad, and was used by Wikipedia until the end of 2001. It continued to be used after this time on special pages such as search results.
	Cunctator logo	2001-2003	This logo was designed by The Cunctator and won the first Wikipedia logo contest, which took place from November to December 2001. The leading submissions of that early contest (which is separate from the much larger international logo contest held in 2003) can be viewed at Logo suggestions/Leading candidates . See #2001 logo contest for more details.

Figure 3.7. Screenshot: Wikipedia's progression of logos (date of screenshot: 9 September 2023).

Wikipedia logo, and frames it as a test logo.

While I was not able to find any conversations about the switch between the American flag and the globe, the implications, in terms of belonging to the community, are obvious. While the initial design was US centric, the globe signals an evolution towards ideals of universality. It is also significant that the first official logo was designed as a result of an open call, opening up design to the community (Wikipedia: Wikipedia logos).

What is more interesting, is the fact that the American flag logo is not a widely acknowledged part of Wikipedia's history. One user mentions the flag in the talk page, lamenting its absence:

Flag

The article content now starts with: "Wikipedia's first true logo was an image that was originally submitted by Bjørn Smestad for a Nupedia logo competition which took place in 2000. Jimmy Wales (Jimbo) thought it would be a much better logo than the flag, and it remained for the next eight months". It briefly mentions 'the flag', but never actually states that 'the flag' (which probably refers to the American flag, which I heard more often being mentioned as the first file that was used as logo) was used as logo. Could someone confirm that this was indeed the case, and can we then make that explicit without keeping it this vague?

(*effeietsanders*)

This user mentions that, at the time of their writing (April 2020), a reference to the flag was present on the main article page. As of March 2023, the reference has been excised. Now the page reads:

Wikipedia's first true logo was an image that was originally submitted by [Bjørn Smestad](#)⁽¹⁾ for a [Nupedia](#) logo competition which took place in 2000.⁽²⁾ It continued to be used after this time on Special Pages, such as search results.

(*Wikipedia: Wikipedia logos*)

The mention of the flag was eliminated in June 2022, after being marked with a "citation needed" tag for years. The edit summary reads "Unsourced", meaning: a reference couldn't be found to support the mention of the American flag. Without postulating a conspiracy to hide the first

Wikipedia logo, I can make some educated guesses as to why the community is not keen to give it prominence: the American flag marks Wikipedia as a US-centric project, which goes against the ethos of universality that animates it.

In doing so, however, it contradicts the feelings of one of the founders, Larry Sanger:

ScottMoonen sensibly suggested:

I'd recommend you change the American flag logo. Extremely ethno-centric et. al. I think a globe logo would be much more fitting, if you want to keep with that metaphor. Or perhaps a book.

Then others responded...

Or a quill and inkwell. -- SunirShah

It's a very temporary thing, the nearest thing Jimbo had to hand. Definitely won't be the Wikipedia logo--not that I agree with you that the U.S. flag is "ethnocentric" (which seems very strange to say to me, given that the U.S. is, after all, the biggest ethnic melting pot in the world). -- Larry Sanger

(*Meta: Historical/Community#OldWikiPediaLogo*)

Sanger, here, is nit-picking on the meaning of the word "ethnocentric" – he is of course right that many ethnicities exist in the US, so the US flag does not represent a specific ethnicity as such. At the same time, though, he is ignoring the subtext in the previous intervention: that the American flag is US-centric. And it does not represent the kind of universality the globe stands for.

The first official logo, according to the Wikipedia logos page⁹, was indeed shaped like a globe. Wrapped around the globe, was the preface to a Lewis Carroll book, *Euclid and his Modern Rivals*, in English:

In one respect this book is an experiment, and may chance to prove a failure: I mean that I have not thought it necessary to maintain throughout the gravity of style which scientific writers usually affect, and which has somehow come to be regarded as an 'inseparable accident' of scientific teaching. I never could quite see the reasonableness of this immemorial law: subjects there are, no doubt, which are in their essence too serious to admit of any lightness of treatment – but I cannot recognise Geometry as one of them. Nevertheless it will, I trust, be found that I have permitted myself a glimpse of the comic side of things only at fitting seasons, when the tired reader might well crave a moment's breathing-space, and not on any occasion where it could endanger the continuity of the line of argument.

(*Wikipedia: Wikipedia logos*)

⁹ As demonstrated by the absence of the flag logo, the Wikipedia page about the history of logos is not an entirely reliable source: it presents a narrative that reflects flatteringly on Wikipedia as a whole. I have decided to use it as a source nonetheless, although carefully, and verifying key claims independently as needed.

The next logo, released on November 2001, was similar, also designed by a volunteer, and also featuring a passage from a book, the *Leviathan*, by Thomas Hobbes, also in English:

Desire to know why, and how, curiosity; such as is in no living creature but man: so that man is distinguished, not only by his reason, but also by this singular passion from other animals; in whom the appetite of food, and other pleasures of sense, by predominance, take away the care of knowing causes; which is a lust of the mind, that by a perseverance of delight in the continual and indefatigable generation of knowledge, exceedeth the short vehemence of any carnal pleasure.

(*Wikipedia: Wikipedia logos*)

At this point in time, English Wikipedia already had company from other Wikipedias. Other Wikipedias adapted the logo to their own local culture in various ways, by translating it, changing its colours, and so forth. The issue of the language finally became pressing, and so an international contest was launched, to find Wikipedia's next logo. A few variants were adopted, until eventually, a logo that was very similar to the current one was adopted.

This logo is much more international than previous logos: it is not wrapped in English language, and it presents characters in alphabets other than the Latin alphabet, namely Armenian, Khmer, Japanese, Greek, Tibetan, Arabic, Devanagari, Chinese, Cyrillic, Korean, Mongolian, Kannada, Hebrew, Thai, and Klingon. The characters chosen are notable for two reasons.

First of all, there were several errors in the logo, documented on their own MediaWiki page (MediaWiki: Errors in the Wikipedia logo): some characters were drawn incorrectly; some versions lacked the correct accents. In terms of aesthetics and belonging, one conclusion that can be drawn is that those characters were not meant to be read by people who actually read the languages they were taken from, and most certainly, no one who spoke the languages represented by the characters had been involved in the design process. Those characters performed universality, without actually embodying it.

The second notable point is the inclusion of Klingon. Klingon is not a human language: it is a constructed language, developed for the television series *Star Trek*, by the fictional Klingon civilisation. *Star Trek* is part of the nerd canon, and within early 00s geek culture, there were some who were actually learning and speaking Klingon as a second language. Essentially, putting a Klingon character on the Wikipedia logo was a sign of allegiance, which both expressed belonging to a group (nerd culture) and was a call for other nerds to join. When read in conjunction with the mistakes made when it came to actual languages spoken by people who were not Western, the message seems clear: the group of Wikipedia editors was defined as geeky Westerners. And this was already an improvement, in terms of inclusivity, on the previous logo, which was clearly steeped in English-speaking culture.

The current logo was introduced in 2010, and presents two changes that are notable from the point of view of the imagined user: the Klingon character has been replaced, and the font of the wordmark (the piece of text under the globe, which reads "Wikipedia – The Free Encyclopedia") was changed. The Klingon character was replaced with a character from an existing, terrestrial

language, Ge'ez (Source: Wikipedia logos). The font change is notable because of the rationale, as outlined in the official Visual identity guidelines. The previous version of the logo used Hoefler, which is a proprietary font. The current logo uses Linux Libertine, an open source font. The stated reason for the switch is as follows:

The puzzle globe typically accompanies the Wikipedia wordmark, originally set in Hoefler, but recently updated to Linux Libertine, an opensource typeface, to facilitate easier localization of the wordmark in new language editions of the project.

(Foundation: Legal: Visual identity guidelines)

Using an open source font means opening up to modifying the wordmark, which was harder to do with a proprietary font. As well as harking back to the FOSS ethos of open source (and indeed, using a font that is named after an iconic open-source operating system), the change practically makes it easier for more people to participate. The boundaries of the Wikipedia community are extended both symbolically and practically.

The history of the Wikipedia logo shows a preoccupation with universality, a value that is implied in Wikipedia's mission, and at the same time is the history of Wikipedians defining what universality is. The use of the American flag at the beginning, combined with Larry Sanger's remark about the multiplicity of ethnic groups present in the US, is a microcosm of the way universality is addressed by Wikipedia as a whole: a small group of people is taken to represent the rest of the world, their internal diversity magnified, and ignoring what they have in common. While it is true that the US comprises of many cultures, it of course doesn't represent the whole world, and most importantly: most people who live in the US are American, and speak English as first, or only, language. So is English Wikipedia: it's run by a relatively small group of people, who are proficient enough in English to navigate the complex process of editing, but write about everyone and everything else. Successive logos are as many negotiations of what "everyone" really means, and the relative importance of various groups of people within the community. Including Klingon instead of a real language sends a clear message: whatever your nationality, if you're into Star Trek, you are one of us. Finally, updating the font to Linux Libertine is both a further nod to the nerd community, and a practical way of working towards a Wikipedia community capacious enough to include anyone.

In sum, design choices related to Wikipedia's appearance – in this case, its interface and logo – embed certain assumptions and desires about who belongs to the community and who doesn't. Such choices do so by purely aesthetic means, through appealing, visually, to certain subgroups of people, and functional means, through affording participation to certain subgroups of people. In the next section, I will, in part, keep examining issues around belonging, by looking at the lineage of Wikipedia's aesthetic, and its roots in modernism. As I will explain below, modernist design was, in some ways, linked to a specific anthropology, and has historical ties to Wikipedia, through Ayn Rand's work, and FOSS ideology.

3.5. Modernism and anxieties around progress

It has been argued elsewhere that there is a broad continuity between modernism and digital or

computer aesthetics (Cram, 2020; Kirkpatrick, 2003; Patton, 2003; Turkle, 1995). In this section, I will focus on how Wikipedia's aesthetic shares two traits with the aesthetic of modernist design: the primacy of function over form, and a tendency to do away with ornamentation. As well as manifesting in stylistic choices and design culture, Wikipedia has historical, ideological ties with modernism, via open software culture and Ayn Rand's work. The free and open source software movement inherited its culture of design sharing from modernist furniture design. Ayn Rand as well as being a modernist writer, was an enthusiast of, and celebrated the ethos of, modernist architecture: *The Fountainhead* is about an architect's heroic defence of modernist principles of architecture, over traditional craft. The underlying ideology that inspires both Wikipedia and modernist design, I argue, responds to shared worries about progress and relevance.

Modernism was a cultural and intellectual movement grown mostly in Europe and the United States, at a time when industrial production of goods was taking hold, from the late 1800s to the mid-1900s (Kuiper, 2024). Modernism cuts across several areas of human endeavour, including art, architecture, literature, social sciences. Here I am concerned with design and architecture. While Ayn Rand's works of fiction may also be considered a product of late modernism (Ashford, 2014), I am not interested here in their qualities as literary works: I will cite them exclusively in relation to design principles.

As a style of design, modernism is dated variably depending on sources and geographical focus; roughly, however, it is safe to say that it was developed, in various forms, between the beginning of the 20th century and 1980. Modernist styles vary, and there are contradictions within the movement itself, but broadly speaking, it was characterised by a rejection of traditional architectural styles, and the ambition to rebuild design on the basis of function (Parsons, 2016). Modernist designers eschewed decoration: the new ideal of beauty relied on materials and functional elements. The engineer was, by some modernists, considered superior to the architect. One of the foundational texts of the modernist movement, by architect Adolf Loos is tellingly titled *Ornament and Crime* (1908/2019). The following quote sums up his position:

We have overcome ornament, we have fought our way to ornamentlessness.
Behold, the time is at hand. Fulfilment awaits us. Soon the streets will gleam like
white walls!

(Loos 1908/2019, p.189)

Simplicity became a marker of progress: shedding ornament meant leaving the past behind, and building a new world on rational grounds. Modernism was not just a style: it was a utopia. It is impossible to disentangle modernist principles from the society that they were inscribed within, and envisaging (Parsons, 2016).

As mentioned above, modernist design was a major influence on Ayn Rand, and the FOSS movement. Ayn Rand applied the principles of rationality and progress to her philosophy, explained both in theoretical essays and works of fiction. The novel *The Fountainhead* outlines her views on society and politics through the medium of architecture (Rand, 1943). Because of the political qualities of modernist architecture, Ayn Rand was able to use it as a shorthand for rationalist ideals of progress and social reform (Schleier, 2002). The main character of *The Fountainhead*, Howard Roark, is an architect, inspired in part by Frank Lloyd Wright, who stands against traditional,

ornamented design, affirming his own individuality against a collectivist society, framed by Ayn Rand as necessarily producing bad architecture (Schleier, 2002). Architecture, in *The Fountainhead*, is talked about in normative terms; it orients the reader within a system of ethics where bad architecture, a tangible mark of moral failure, is produced by a society stuck in its ways, held back, in part, by the reliance of its members on each other. Good architecture, on the other hand, is the result of the work of a lone genius – Roark – affirming his independent thought against widespread conformist pressure. Rand’s *Atlas Shrugged*, while not focused on architecture, similarly embodies a spirit of social reform (Younkins, 2017) akin to a modernist utopia. I am not going to examine the fine detail of how Rand’s thought and modernism interact (for an analysis, see e.g. Schleier, 2002 and Ashford, 2014). It is sufficient for me to show that such interactions exist, and have influenced Wikipedia’s aesthetic and ethos.

The other channel through which Wikipedia has inherited elements of modernism is the open source software movement, inspired by the modernist innovation of creating freely reusable blueprints for furniture. The emphasis on function and simplicity made modernist furniture easily reproducible, in factories and workshops. Openness is a fixture of modernism since the beginning:

Any joiner can make the objects that I use in my interior design. I am not making myself a patent architect. Any marbler, any textile maker or industrialist can make my things and doesn’t need to ask me respectfully. The main thing is that he does an honest job.

(Loos 1924/2019, p. 266)

Similarly, in the 1970s, Enzo Mari’s *autoprogettazione* format allowed anyone to build their own furniture by using blueprints, shared freely by Mari himself and others (Mari, 2002). The ethical qualities of modernist aesthetic are upheld by a firm commitment to minimalism, a proxy for efficiency and reusability.

The conflict between ornament and minimalism didn’t end with furniture. A similar, and similarly morally charged, dynamic occurred in the history of Web design. Early 1990s Web design was characterised by an explosion of creativity: websites were colourful, busy, and, often, animated (Ankerson, 2018; Ford, 2019). The dot-com boom favoured an abundant, maximalist style, thanks to innovations such as Flash and CSS (Ankerson, 2018). With the advent of Web 2.0 in the early 2000s, anyone with access to an Internet connection became able to create their own space on the Web even without coding skills, thanks to tools such as Dreamweaver, which allowed to create a website without coding, and platforms like MySpace and WordPress, where users could host a personal page and customise it to their taste. Users of no-coding Web design platforms took advantage of their online presence to express themselves, through skins, fonts and multimedia content. Busy websites and colourful designs started to be associated with amateur designers.

This maximalist approach to Web design was so popular, that those giving Web design advice had to beg designers to keep minimal. The following quote is from a 2002 article about the possibilities of “hypertext”:

Adopt one [font] for your body text and a second font for your headlines. Use both consistently [...] Use emphasis devices such as bold or italics sparingly. [...]

Be careful with your use of color on the Web. It is a very powerful emphasis device, but you run the risk of making your pages look like Mardi Gras floats and alienating people who are color-blind.

(*Felker, 2002, p. 335*)

As the experience of users started to become important in the eyes of Web designers, functional websites became the gold standard of Web design (Ankerson, 2018). Over time, a more pared-back Web aesthetic started to win out, coloured by ethical concerns: overly decorated websites, as explained in section 3.2.1, fared poorly in terms of accessibility. Apple.com had a pioneering role, with its 1998 redesign, that reflected the minimalism of Apple products and, significantly, used a font, Hoefler, that was later adopted for the Wikipedia logo. The shift between exuberant and subdued designs put function before form in computational terms as well: heavily decorated websites, as well as being difficult to use, required a lot of computing power, and were often slow to load, because of the large amount of data needed to display elaborate graphics (Ankerson, 2018), echoing the modernist association between a light aesthetic and efficiency.

As well as responding to concerns around accessibility, minimalist Web design was a sign of belonging to the community of professional Web designers. First, because it showed that the designer understood, and cared about, industry standards. Secondly, because it showed a commitment to craft. Kennedy (2012) reports how, as Web design was becoming professionalised, Web designers would lament that tools such as Dreamweaver and Flash were degrading their craft, by putting a barrier between code and designers. The code powering websites created with Dreamweaver was rather opaque: it was written as a set of instructions for a machine, not for human eyes. As a result, it was difficult for coders to make any changes (Kennedy, 2012). Flash fared even worse, as some parts of Flash websites were completely obscure: their code couldn't be accessed or modified at all. The inability to intervene threatened the craft of coding by stifling learning, in a context where most coders learnt through practice, and often started out by looking at, and modifying, other people's work; and by preventing coders from upholding ethical standards by improving the design of ready-made websites (*ibid.*). The discomfort felt by designers at the time, I would add, might have been bolstered by the disregard for the social qualities of coding – the glue that holds programming culture together (see 2.3.3.3 for context). Opaque code is also, radically, anti-modernist and anti-openness: it hides function and hampers reuse.

The battle against decorative Web design is relevant to Wikipedia because of the values it promotes, and its role in marking belonging to a specific community. I have written above how minimalism, in furniture design, expressed a commitment to efficiency and reusability, by making it easier for furniture to be built in a factory, and blueprints to be used by amateurs. Minimalism, in a modernist context, also embeds two values that are central to Wikipedia: civility, and craftiness.

The first value is something that Wikipedians call “civility”, harking back to a distinction from the constructed category of savagery. The Faulkner quote, above, argues for a simple aesthetic by evoking accessibility, and a more general disdain for maximalism, expressed through the reference to Mardi Gras floats. The Mardi Gras reference evokes chaos, disorder, and a break from the productive, everyday mode of existence. Compare with this passage from Loos:

The Papuan slaughters his enemies and devours them. He is not a criminal. But if modern man slaughters and devours someone, he is a criminal or a degenerate. The Papuan tattoos his skin, his boat, his rudder, in short everything that lies to hand. There are prisons in which 80 per cent of the inmates have tattoos. The tattooed people who are not in jail are latent criminals or degenerate aristocrats.

The urge to ornament one's face and everything that lies to hand are the primal origins of visual art. It is the babbling of painting.

(p.188)

In the same essay, Loos explains how the intellectual superiority of “modern” (Western, of course) “man” means that “he” has overcome ornament. This theory is placed within an evolutionary framework, evoked also to the reference to degeneration theory, popular at the time as the idea that evolution can be reversed, and that when this happens, inferior people appear. Degeneration theory placed within the group of degenerates everyone who was not productive in a capitalist sense, chiefly women, children, criminals and aristocrats (Pick, 1989), and has had a strong influence on aesthetics in modern times (Dijkstra, 1986). The “Mardi Gras” reference above has a similar cultural quality: it evokes moral debauchery, lassitude, and a break with the rules of ordinary life, which regulate the rhythms of work. The association between degeneration and decoration is also present in Ayn Rand, following the theories of architect Louis Sullivan, who explicitly saw modernist design as an antidote to the degeneration of American – presumably, white – masculinity (Schleier, 2002). Minimalism is, then, a proxy of something we might call “civility” or “professionalism”.

Within an American context, Mardi Gras is also associated with Louisiana, and is important to Black culture. It is, in other words, a group-specific, situated practice. Minimalism, on the other hand, is often framed as universal: in absence of decoration, an object's function, or in other words, its essence, takes centre stage – making it easily shareable, translatable. In his essay “On Thrift”, Loos praises minimalist objects for fitting next to any other object, from anywhere in the world; they embody perfection; they are timeless, as people are less likely to become tired of them.

Wikipedia's minimalism also allows direct manipulation of markup or code, thereby encouraging editors to learn technical skills, and turning them into members of a craft community. Wikipedia is a brutalist website: it wears its materials on its sleeve; the bare bones of the website are readily accessible – the source editor shows markup, stripping off a layer of interface to show, at least to some extent, the workings of the website. Turkle (1995) noted how early personal computers, by requiring the user to type simple commands, followed a modernist ethos; they encouraged engagement with mechanisms, fostering a sense of participation, of belonging to the community of makers. The graphic interface, featuring a desktop modelled on real objects, represents a shift towards the postmodern (Turkle, 1995). Wikipedia's technical editing interface and unfinished look, inviting users to participate (as argued in section 3.2), echo the same sensibility. As I will argue in chapter 4, the relationship between coding, craft communities, and the transmission of knowledge is essential to understand Wikipedia.

Wikipedia is modernist in its ambition to universality, in encouraging editors to act like professionals, and in placing editors in a craft tradition. Minimalism, has an epistemic role as well: it

validates knowledge by performing trust by proxy: it is a way to show that the website is made by people who know what they are doing and are reliable witnesses.

Both modernism and Wikipedia respond to anxieties about progress and relevance. The modernist movement took hold as the category of the new emerged, and with it, the ideal of novelty (Levenson, 2011). Modernist designers and architects were concerned with capturing the spirit of the times, as opposed to following tradition (Parsons, 2016). Similarly, conversations around Wikipedia's design hinge around the idea of keeping up with the times. Some examples:

Dear all, Today I learned that, despite having \$100 million in the Endowment fund, we can't have a design team big enough to make our websites not look like they're stuck in 2001. I don't know if anyone is behind the wheel, but the car is expensive. Sincerely, Galder

(Galder Gonzalez Larrañaga)

Indeed, the lack or modernisation of the Web interface, and lack of an improved Android/iPhone (or simply "smartphone") editing app[1] are possibly some of the major areas to focus. If the "next billion Internet users", a term we often used to use are not getting involved as much as we expected, possibly "interface" is one reason behind it. [1] I am aware of the currently available apps.

(সুজ, ১৫ আক্টো.)

Similar worries are reflected in Wikimedia Foundation's white papers:

By understanding our users needs and expectations we can modernize our products, and provide a user experience that informs and delights.

(MediaWiki: Experience [white paper], p. 30)

Connected to the concerns around modernity are worries around relevance, and the number of editors: if Wikipedia becomes obsolete, the fear goes, then people will not read it, which will bring down the conversion rate of readers into writers. And for the project to work, large numbers of editors are needed:

On the contributors' side, Wikimedia projects compete with modern platforms that provide gentler on-boarding and guidance to new users. The competing platforms provide rich, multimedia editing tools and emotionally reward their users with explicit gratitude, meaning, and status.

(MediaWiki: Experience [white paper], p. 29)

In 2001, when Wikipedia was launched, a modern look meant following the aesthetic canons described above. In order to achieve a contemporary feel, designers had to get rid of the vestiges of the dot-com boom aesthetic: the bright colours, heavy-to-load flash animations, busy interfaces and so forth. It is unclear, to me, that Wikipedia actually ever did look modern – concerns with looking obsolete have been present from the beginning. The signifiers of modernity, in the Wikipedian discourse around appearance, have also changed over time. At first, as described above, complaints

concerned the website looking too technical, and as such, being impenetrable to outsiders (section 3.4). Nowadays, the worries are mostly around the content being boring – hence the centrality of multimedia in discourse.

Wikipedia discourse also presents a strain of nostalgia and conservatism, which reinforces my claims about a mainstream tendency to chase progress. A practical example of this phenomenon can be found in the range of skins available to Wikipedians, and the early discourse around skins themselves. First of all, even at the very beginning was already present in conversations a pushback against new skins. The following quotes are from a mailing list conversation dated 2004:

Personally, I don't consider [designing a new skin] a development priority, and it's unlikely that it will happen anytime soon.

(Erik Moeller)

Well, you may have noticed that even minor changes to the skin meet with quite strong resistance and have to be defended individually. Realistically, you will have to defend your new skin, and it will probably end up being voted on, before it is made the new default.

(Erik Moeller)

The skins that are currently available have been developed over time. Some are older than others – the old ones are kept to appease nostalgic editors. When new skins were developed, and the issue of changing the name of old skins came about, names such as “Nostalgia” and “Gold standard” were suggested for older skins. The presence of an attachment to older versions of the Wikipedia interface can be detected in the appearance of the skins themselves, as well. Wikipedia, at the time of writing (March 2023) offers its registered users the option to pick their skin among 5 available ones. While there are some differences, some of the skins are almost identical: given the attachment of certain members of the community to old designs, it seems fair to conjecture that the almost-identical designs might have been preserved to avoid backlash.

In conclusion, the design of en.wikipedia.org carries a distinct modernist quality: en.wikipedia.org is a minimalist website, placing more importance on highlighting technical features, useful for editing, over a pleasant appearance. In doing so, it addresses an underlying anxiety around progress and remaining relevant, eliciting, in turn, a conservative reaction from some pockets of the community. In line with the ideology of some modernist architects – namely, those who had an influence on Ayn Rand, herself an inspirational figure for Wikipedia – Wikipedia's minimalism is connected to intellectual virtues, and aimed at projecting trust, by embedding an ideal of professionalism encoded in early 2000 Web design advice, performing an identity connected with intellectual endeavour, and distancing itself from inefficient uses of technology. This last aspect of Wikipedia's minimalism, harking towards elements of identity as markers of epistemic superiority, is, as I will argue in the next section, connected to the figure of the nerd.

3.6. Geek culture, antiaesthetic and antisensuality

Wikipedia's aesthetic hides human labour, while leaving infrastructure bare. I will suggest in this

section that this double movement of concealment and baring is the expression of what Sherry Turkle calls an “antisensual aesthetic” (1984), prevalent in nerd culture. Thus Wikipedia’s aesthetic acts, I argue, as a form of boundary maintenance. The emphasis of intellectual, over bodily, pleasure also frames knowledge as abstract, reinforcing the Wikipedian doctrine of neutrality as neutral point of view.

This point constitutes an anchor for a wider trajectory: hiding the bodies of editors, as I will argue over the course of this thesis, informs, at a deep level, the way in which Wikipedia’s knowledge-production process is organised, and has seeped into large parts of its design. By the same token, as I will show in section 4.5, challenges to Wikipedia’s core values have taken the form of showing bodies, by making bodies visible, hearable, or perceivable, or even just hinting at their existence, through correlates such as affect and identity.

Sherry Turkle, in *The Second Self* (1984), talking about the artistic preferences of geeks, notes how the way one felt when faced with a work of art was not important in nerd culture at the time. Rather, the emphasis was on the structure of works of art. By way of illustration, she points towards the cult book *Gödel, Escher, Bach*, well-known and revered in nerd culture, which focuses on the structural continuities and dialogue between Gödel’s theorem, Escher’s art and Bach’s music. Anecdotally, she points towards one of her participants disparaging the abstract art in her study, and bringing her an Escher print instead, because “it gives you things to think about” (p. 202). The same participant appreciated Bach because of the structure of his fugues – to the point that the quality of sound seemed irrelevant: Turkle reports him listening to music from a scratched record, that made a rather unpleasant sound. This disregard for sensory pleasure, according to Turkle, goes hand in hand with a more general disregard for the body, which is disciplined and bent out of shape in order to serve the machine, nurturing the stereotype of the the nerd as “ugly” (Turkle, 1984).

Turkle’s account is consistent with other descriptions of nerd aesthetic as, in essence, an antiaesthetic: a performative, proud display of shabbiness and disregard for appearances, in terms of personal grooming (Ensmenger, 2015) and living spaces (Kendall, 1999). Turkle connects nerd antiaesthetic with a form of antisensuality:

The sensual goes beyond the overtly sexual. There is sensuality in music, literature, art, and there are sensual relationships with the world of things: the musician caresses his or her instrument; its shape, its tonality, its touch, can be pleasing and exciting. But the prototypical hacker’s taste in each of these realms tends not toward a sensual caress but toward an intellectual contact.

(Turkle, 1984, p. 201)

In other words, nerd aesthetic is based on a bracketing, or even mortification of the body, by ignoring one’s needs, forgoing personal care, and subduing one’s routine and lifestyle to the needs of computers. It has been noted elsewhere how digital technologies enable the bracketing of the body, generating the temporary fantasy of not being embodied at all (Boler, 2007; Buongiorno, 2019).

Wikipedia’s cultural understanding of the self is, indeed, antisensual: Wikipedians are not supposed to have a point of view, nor bring in affect or identity to their work on the encyclopaedia (see Appendix A for details, and chapter 4 for the implications of this framing). As acknowledged in the quote above, Wikipedians don’t receive emotional rewards when they edit; the experience of

editing is geared around intellectual gratification, as opposed to pleasure. When it comes to designing digital infrastructure, antisensuality manifests as the tendency to abstraction described in section 3.3.1. Fleshless text is preferable to more sensorily rich media, such as, for instance, video. The pursuit of platform-independence further extends the ideal of disembodiment, by rendering irrelevant the material substrate on which information is carried, arching towards an ideal of knowledge as pure information.

A tendency towards antisensuality in terms of bodily enjoyment, however, doesn't translate into a complete abstinence from aesthetic pleasure; not within geek culture, nor on Wikipedia. Rather, aesthetic enjoyment, in accordance with modernist ideals, is found in function. In hacker culture, aesthetic pleasure is often connected with the beauty of code. Beautiful code is, mostly, minimalist code: code that does what it needs to do in the least possible number of lines (Ensmenger, 2015). Beauty in this context, then, consists in something that we might otherwise call efficiency, occupying an analogous space to the idea of elegance in scientific theories. Geek culture also presents canonical forms of entertainment, such as making amusing code (Ensmenger; others) and indulging in specific pastimes, like the role-playing game *Dungeons and Dragons*. Pleasure, however, as anyone who has ever played *Dungeons and Dragons* knows painfully well, is highly regimented, subject to the outcome of calculations, or enjoyed as part of an intellectual or technical pursuit.

Steven Levy, in a rather celebratory book about early hacker culture, describes the intermingling of pleasure and efficiency as follows:

A certain esthetic of programming style had emerged. Because of the limited memory space of the TX-0 [an early computer, used throughout the 50s and 60s] (a handicap that extended to all computers of that era), hackers came to deeply appreciate innovative techniques which allowed programs to do complicated tasks with very few instructions. The shorter a program was, the more space you had left for other programs, and the faster a program ran. Sometimes when you didn't need speed or space much, and you weren't thinking about art and beauty, you'd hack together an ugly program, attacking the problem with "brute force" methods. "Well, we can do this by adding twenty numbers," Samson might say to himself, "and it's quicker to write instructions to do that than to think out a loop in the beginning and the end to do the same job in seven or eight instructions." But the latter program might be admired by fellow hackers, and some programs were bummed to the fewest lines so artfully that the author's peers would look at it and almost melt with awe.

(Levy, 1984, p.36)

"Bumming" in this context refers to the practice of pruning code in order to make it more efficient. The language of affect and beauty ("melt with awe"), here, is mixed with technical considerations (efficiency is prized because it makes good use of the limited available memory space) and social acceptance. Competitive coding – that is, designing code that is more beautiful, or does a specific, even if useless, thing – and hacking – breaching systems for the sake of it, without any material gain – have been features of geek culture since the early days, and still are.

Two things are notable here: the fact that code itself is beautiful – so, while the body can't be

beautiful, the machine can – and that there is some kind of pleasure associated with creating beautiful code. On the first point, it is helpful to note that technical beauty is an important element of nerd culture. It has been noted, for instance, how game consoles have evolved over time to show, and show off, the workings of the machine: due to the necessity of adapting existing hardware to a fast-evolving commercial landscape, the practice of modifying one’s equipment became widespread, and a game in and of itself (Simon, 2007). Hence the machine itself constituted a point of admiration, and its parts – e.g. lights, fans – became aestheticized; bare internal circuits, initially exposed to prevent overheating, have become a stylistic fixture (*ibid.*). By the same token, iMac computers were built with a transparent, but colourful, casing, which both showed and beautified the internal mechanisms that made the machine work.

Analogously, Wikipedia, while hiding the bodies of editors (as shown in 3.3.3.1) displays the material means by which knowledge-production is achieved. “Edit” buttons and “citation needed” tags, which make it possible for content to be updated and extended, are visible to everyone – contrast with a possible alternative interface where only logged-in users can see them. The default editing interface is made of Wikitext, the markup language carrying instructions about how text, images, and other media are displayed on the reading view. Community, policy, and other administrative pages are not tucked away, but displayed on the left. Giving direct access to editing infrastructure is helpful in terms of recruitment, as explained in section 3.2. But displaying the inner workings of the encyclopaedia is not just a practical choice – it’s an aesthetic one. Wikipedia doesn’t just put function over form: it shows off function. As articulated above, its bareness is a point of pride, expressing an ideal of precision over frivolity, rationalism over irrationality, professionalism over amateurism. Wikipedia presents itself as a working encyclopaedia – as a workshop, not a window display.

Wikipedians tend to be quite coy about the pleasure of editing: when asked why they edit, my interviewees typically appealed to a sense of duty, or as a means of learning new things. This is consistent with surveys, where Wikipedians have indicated reasons for editing related to learning (Simon, 2007) or contributing to a community (Lai & Yang, 2014). Historically, user pleasure has been downplayed by Wikipedia’s designers, as illustrated in the conversations among the initial group of designers (see section 3.4). Strategy documents take great pains in arguing for the importance of entertainment value, which is framed as a departure from Wikipedia’s current style. I have already remarked on Wikipedia’s house style’s programmatic, ideologically motivated dullness (section 3.3.1).

Enjoyment is culturally accepted only in the form of the satisfaction of an intellectual pursuit, analogously to the awe of a hacker in the presence of efficient coding, or their appreciation of structure in art. By the same token, the Wikipedian response to the problem of dullness has been, historically and currently, technological, rather than aesthetic. When discussing the length of articles, for instance, early Wikipedians and current Wikimedia Foundation staff alike recognise that long, strictly factual articles don’t make for riveting reads. However, the idea of adopting a more engaging style is never contemplated: the antisensual choice is to find a technical fix that appeals to structure, rather than one that appeals to pleasure. Consequently, interventions are suggested around making articles shorter, or packaging information in multiple formats, where video is assumed to be more engaging than the written word, by virtue of containing more information, and in any case, as a concession to those who are not capable of finding enjoyment in “pure” text.

Assumptions around what is and is not interesting or aesthetically appealing are also entangled with ideas of belonging: as shown throughout this chapter, a certain condescension permeates discussions on Wikipedia's style and necessity to appeal to a wider public, where the wider public is often framed as less educated and intelligent than the speakers – "lusers", as opposed to current Wikipedia readers, which are capable of "deep dives" and engaging with "challenging content".

In this sense, and unsurprisingly, Wikipedia's antisensual aesthetic is a re-affirmation of geek identity, which of course, ironically, situates it very firmly within a specific culture, and points towards a specific embodiment. Interestingly, Wikipedia's design achieves the opposite of what it seeks: in attempting to perform ideals of inclusiveness and universality, it ends up reinforcing existing barriers to participation, and retraces the epistemic landscape that was drawn when Wikipedia was founded.

Conclusion

In this chapter, I have analysed the main influences on Wikipedia's design, and explored how foundational Wikipedia values are embedded in Wikipedia's aesthetic.

Propositional knowledge has a multifaceted role in Wikipedia's design, through its modularity and, hence, its shareability. On one hand, it lends itself to be accessed in via hyperlinks, and consumed through cultural practices that are fundamental to Wikipedia's readers (the social fact-check and the rabbit hole, section 3.2). On the other hand, shareability allows abstraction from provenance, and for knowledge to be consumed and distributed in a manner that detaches it from its source. Indeed, modularity allows for the work of editors themselves, who can contribute any size of information to the page, ranging from fixing a typo to writing a whole section. Propositional knowledge allows radical collaboration, also a fundamental principle of Wikipedia's ethos.

Wikipedia's aesthetic can be described, broadly speaking, as minimalist. From the subdued tone of voice to the simple appearance of its interfaces, Wikipedia follows a Web design trend that equates simplicity with professionalism and trustworthiness, following a time-honoured tradition that started with modernist design in the early twentieth century. Ornament is added begrudgingly, as a necessity for survival, and implicitly framed as a tainting of Wikipedia's core ethos. A tension emerges between abstraction and situatedness, that is never quite resolved, and spills over into proposals for reform of Wikipedia's interface. As I will argue in chapter 4, the same tension is present in how the knowing subject, rooted in a Randian understanding of neutrality, is constructed through technical infrastructure.

Wikipedia's aesthetic is intertwined with its epistemology, in two ways: by expressing epistemic tenets developed through policy and community practice, and by constituting part of a system of validation that directly lends credence to the content of the encyclopaedia. Thus, abstraction can be read as expression of neutrality, and shedding ornament, in favour of highlighting ways of interacting with the encyclopaedia, as a way of framing Wikipedia as unfinished, implicitly embedding a conception of knowledge as constantly evolving. Further, minimalism projects an ideal of professionalism aimed at eliciting trust, interconnectedness produces a sense of cogency, and uniformity conceals the fragmentary nature of Wikipedia's knowledge-production process, signalling a universal, neutral point of view.

It is worth reminding here that Wikipedia's appearance has the difficult task of conveying both accessibility and trustworthiness: users, as readers, need to feel both like they can contribute, and like they can trust others who contribute. This sense of trustworthiness is partly achieved through a minimalist, professional appearance, partly through signalling, via aesthetic means, that the website is made for and by people like the reader. Wikipedia's aesthetic – its interface, language, image policy – is geared around structure and function, rather than decoration, communicating that Wikipedia is about substance, not appearance – which is what people want in a reference work. Secondly, a functionalist design aligns with aesthetic values that are common in nerd communities, thereby defining the community it is made by and for; it is a form of boundary maintenance. Privileging nerds also means appealing to an audience that is deemed likely to want to contribute to an encyclopaedia, since, as I will articulate more precisely in the next chapter, the figure of the nerd is defined by an affinity with knowledge consumption and production.

Pure aesthetics is, in this sense, just the tip of the iceberg: in the next chapter, I will expand on how Wikipedia's design contributes to boundary maintenance, and on the fundamental role that programming plays in Wikipedia's epistemology. By showing the connection between programming practice, culture, and Wikipedia's infrastructure, I will lay the groundwork for my theoretical contribution: in the final chapter, I will argue for the necessity of conceptualising the design space as a relational space where the subjectivities of designers and the materiality of tools and artefacts contribute to define values.

4. Programming an encyclopaedia

222. ENCYCLOPEDISTICS. Not the essential—characterizes—not the main bulk—rather the inessential—the peculiar.

Novalis, Notes for a Romantic Encyclopaedia

4.1. Introduction

In this chapter, I argue that the practices of editing Wikipedia’s content and coding the infrastructure that sustains it are deeply intertwined: knowledge production on Wikipedia is not just mediated by technology, it’s carried out, to a large extent, through technology, by coding artefacts that produce, organise, and filter knowledge. Conversely, in my view, designing technology for knowledge production is an epistemic activity in and of itself: epistemic values are embedded and upheld through coding, with important consequences on the encyclopaedia as a whole. Wikipedia’s epistemology is also, in my view, heavily influenced by coding practice: much like its style is rooted in nerd aesthetic, its core values, including epistemic values, in part come from, and in part are defined through, coding.

Through a series of case studies, I will show how specific coded artefacts that form part of, or are active on, Wikipedia’s website, carry out epistemic labour by operating distinctions between contributions and contributors. As a result, these artefacts embed and uphold specific, sometimes conflicting, epistemic values, thereby becoming a space for the development of Wikipedia’s epistemology. I will put my observations into context, by drawing connections between coding on Wikipedia and software development culture and practice. I will conclude by framing knowledge-production on Wikipedia as a technoepistemic practice, which I will define as a practice that blends knowledge-production and design, based on a regulating principle that I name poietic objectivity.

4.2. Coding as means of knowledge-production

In this section, I explore the role of coding, and coded artefacts, in the process of knowledge-production on Wikipedia. As well as being an epistemic community, Wikipedia is a coding community. As outlined in chapter 1, one can think of Wikipedia’s coded infrastructure as being made up of a bedrock of core code – maintained by Wikimedia Foundation staff and amateur engineers coordinated through a platform called Phabricator, which tracks tasks related to various projects – and bespoke code (Geiger, 2014), the ever-growing suite of small programs, such as bots and gadgets, created by Wikipedians to support their editing work. Many artefacts that are now an integral part of the Wikipedia technical ecosystem were created by volunteers. An example above all: the software that powers Wikipedia, MediaWiki, was created by an early volunteer, Magnus Manske, who is celebrated yearly on Magnus Manske Day (Wikipedia: Magnus Manske Day).

Coding practice, on Wikipedia, features at every stage of knowledge-production: writing articles, organising content, and validating its accuracy. Saying that coding pervades the whole process of content production doesn't mean saying that all Wikipedia editors are coders; rather, that some of them have to be: without editors who are also coders, Wikipedia couldn't exist in its current form. If we frame Wikipedia's authorship as collective – and I argue we should – then coding is part and parcel of the process of knowledge-production, as, in order for Wikipedia to function, it is necessary that at least part of the community engages in coding practices.

4.2.1. Writing as, and, through coding

In this section, I will show how the practices of writing and coding are intertwined in Wikipedian practice, by examining how editing is performed using the default editing interface and automated through the coding of bots.

Wikipedia currently offers two editing interfaces: the Visual Editor and the Source Editor. The Visual Editor is a WYSIWYG (What You See Is What You Get) interface, that looks very similar to the article page, and can be modified by pointing and clicking. Wikipedia's default – and oldest – editing interface, however, is the Source Editor. The Source Editor displays a markup language called wikitext, native to Wikipedia, and similar, but not identical, to HTML. In order to edit the content of an article, human editors write or modify both the article's text, and the surrounding markup, which specifies how content will be displayed. Depending on the definition of "coding" one adopts, modifying markup could, in and of itself, constitute coding practice; if that were the case, then many Source Editor users are both writing and coding at the same time¹⁰.

Even taking the view that wikitext is not a programming language, it is important to note that markup and code share a fundamental quality: they are both kinds of text that does things. As Alexander Galloway observed, "*Code is the only language that is executable*" (2004, p. 165, original italics). In other words, code is text that does things, or as Galloway once again puts it, code is "materially affective" (Galloway, 2004, p. 244). Whether markup is code or not, it is, undoubtably, text that does things: editing or writing in a markup language has consequences that reach beyond changing the string of letters and symbols itself. In the case of wikitext, markup can introduce a reference, change the formatting, introduce links between pages, and so on. Wikitext, at the same time, sits physically, on the screen, alongside inert text – Wikipedia's content. Changing inert text doesn't do anything other than changing the text. The continuity between active and inert text has consequences in terms of practice and ethics, which I will explore in section 4.3.

Wikipedians also use more traditional forms of coding to produce and edit content, by creating and deploying digital artefacts known as bots. Bots are small computer programs designed to perform specific functions on Wikipedia's content. Some bots are capable of writing whole articles, by finding open access sources on the Web and importing information into Wikipedia. This operation includes

¹⁰ Whether HTML – and by extension, all markup languages, including wikitext – is a programming language is a matter of dispute. Consequently, whether modifying or writing HTML can be classed as coding, depends on whose side one takes. I will give more details about this dispute in section 3. In any case, nothing can be gained, in terms of my analysis, from attempting to settle the matter. Rather, I will simply state that markup, at the very least, shares some characteristics with code, which are relevant to present purposes, and that the practice of editing markup can be, in certain ways that are significant in terms of the process of knowledge-production on Wikipedia, similar to coding.

“wikification”, i.e. transforming external data and text into Wikipedia content via a set of standard operations such as, for instance: turning HTML into wikitext, linking new content to pre-existing Wikipedia entries, fitting text into a template, and categorising articles based on Wikipedia’s system of categorisation. The first examples of articles created with the use of bots date back to the beginning of Wikipedia: “several hundred” articles were created in February 2002 by selecting and wikifying entries of an American telecommunications glossary, the Federal Standard 1037C; ClueBot II created articles on asteroids based on NASA data; entries from the 1911 edition of Encyclopaedia Britannica have also been imported (Wikipedia: History of Wikipedia bots). A bot called Rambot has created more than 30,000 entries about small towns, based on data from the CIA World Factbook and the US census (Niederer and Dijck, 2010). Articles created by bots can of course be further edited by humans. Sometimes human intervention is necessary: when hundreds of entries were adapted from the Easton’s Bible Dictionary, a reference book published in 1894, the resulting articles were criticised for their angle and language, and had to be tweaked manually (Wikipedia: History of Wikipedia bots).

Bots are extensions of their makers – both as tools, and because their user profiles (bots are formally editors – in some cases, with more extensive powers than many human editors) are connected to their designers’. Bots are described, on their profiles, with the name of their operator, and a link to their user page. Bots share their achievements with their creators: the edit count of an editor using a bot includes the edits made by the bot (Wikipedia: List of Wikipedians by number of edits). The tight coupling of bots and their operators is often reflected in their names: Rambot, mentioned earlier, is named after its designer, user Ram-Man (his real name is Derek Ramsay). The way Ramsay talks about Rambot is also significant; in his user page, he writes: “Here I am most well known for having authored thousands of articles on U.S. cities and counties using the rambot and other automatic article generation techniques” (User: Ram-Man). In Ramsay’s description, Rambot is not a co-author: it’s a tool Ramsay himself uses to write articles. Rambot is no more a co-author than a laptop or a pen. The practice of coding then, is framed as part of the practice of writing: coding Rambot was one way in which Ramsay wrote content.

In sum, the practice of editing Wikipedia blurs the lines between writing and coding, to various degrees. In the case of editing markup, while Wikipedians may not be programming under certain definitions of programming, they clearly manipulate text that enacts certain effects: the practice of editing markup is at least, as a practice, adjacent to that of coding. Even when editors don’t edit markup at all, they are editing text that is physically close to markup. As I will argue in section 4.3, this continuity has consequences in terms of practice, because it gives editors the feeling of manipulating a special kind of text, which comes with a sense of responsibility and, in some cases, worries around one’s skills. I will also argue that the material adjacency of markup and inert text is significant in ethical terms. Coding bots can be more straightforwardly framed as a component of Wikipedia editing: a bot is a coded artefact, created to extend the writing practice of its designer. Coding and writing practices, then, are intertwined in Wikipedia’s context. In the next section, I will describe how coded objects participate in Wikipedia’s process of validation of knowledge.

4.2.2. Validation through code

Every system of knowledge-production features a method of certification of knowledge, aimed

at sorting truth from falsehoods; that is: a system of validation. Any trust placed in a knowledge-production enterprise appeals, to an extent, to its system of validation: for instance, we believe in scientific truth, to the extent we do, because we trust the scientific method and the vetting of results performed by peer-review systems. Naturally, Wikipedia has its own validation system, which is, I will argue in this section, technical. In my view, Wikipedia's coded infrastructure functions as a system of validation by managing the sourcing and manipulation of information.

Since Wikipedia can be edited by anyone, the expertise of editors cannot function as a guarantee of truth. Rather, the accuracy of Wikipedia's content is guaranteed by the sources of information relied upon by editors: that is, provenance, enacted through a technical system that traces the lineage of Wikipedia's content. Such infrastructure is composed of hyperlinks, tags, and bots tending to the longevity of hyperlinks.

Each statement within a Wikipedia article should, ideally, be supported by a citation, or in material terms: a link to the source of the information provided. Because of their role in citation practices, hyperlinks are essential to the integrity of Wikipedia's knowledge-production project: they are the direct line of communication between original research, performed outside of the Wikipedia community, and the encyclopaedia itself. As well as guaranteeing provenance, the pervasiveness of hyperlinks frames Wikipedia's content as secondary literature: citations signal to the reader that nothing on Wikipedia is new – rather, that Wikipedia is a mere aggregate of information found elsewhere, consistently with the core policy of “verifiability, not truth” (Wikipedia: Verifiability).

When a source is not available, a tag that reads “citation needed” is added at the end of the statement. The citation tag has become a somewhat iconic feature of Wikipedia, embodying its commitment to accuracy by mimicking academic practices. The citation needed tag performs multiple roles, depending on audiences. A Wikipedia reader, who is not a contributor, will interpret the citation needed tag as a warning that the information provided might not be accurate, which in turn will inform their behaviour: since the statement in question is not fully substantiated, they might want to take it with a pinch of salt, or seek clarification elsewhere. By implying that, ideally, all statements should be corroborated, the tag also highlights how important accuracy is to Wikipedia: admitting a limitation can foster trust. A contributor, on the other hand, sees the “citation needed” tag as a call to action: it invites them to look for a source. The tag can also turn a reader into a contributor, who might be enticed into performing the seemingly simple task of adding a reference.

As important as they are, hyperlinks are, by their nature, unstable: they connect Web pages, which may be modified, deleted, moved. For hyperlinks to function as a structure of validation, then, links need a support system. The phenomenon by which links remain but their destinations no longer exist is known as “link rot”; untethered links themselves are known as “dead links”. Link rot is addressed on Wikipedia through a set of bots performing layered automatic archiving: when a link is added manually to a page, the page where the link lands is automatically saved on an external archive, as a backup. The bot WaybackMedic (now in its 2.5 version) checks for link rot and, if a link is dead, it resurrects it by replacing it with a link to the corresponding archived page.

To summarise this section so far: provenance, on Wikipedia, is maintained through a dynamic system of coded artefacts. Wikipedia's content is scaffolded by an ever growing system of hyperlinks, protected by ancillary technologies that encourage growth (tags), and preserve what has already been achieved (bots and archiving infrastructure).

Once information has been collected, a system needs to be in place to ensure it is handled appropriately: since Wikipedia is edited by a large number of non-experts, information is managed through a process that doesn't place excessive weight on individual contributions, and contains a built-in error-detection system. Given the kinship between Wikipedia's founders and programming culture, the system in question is also strongly inspired to an existing method of managing collaboration on code development, namely the process known as "versioning".

The simplest form of versioning is making subsequent, improved drafts of the same text, while systematically archiving old ones. Writers creating increasingly more refined drafts are, mostly unknowingly, using a versioning method. Versioning in programming is more complex, because it is designed to allow for decentralised cooperation: a group of developers work on the same piece of software by making copies of it, modifying them, then committing changes to the main body of code. Versioning, in the form of CVS, or Concurrent Versions System, was devised to respond to coordination needs in small groups of software developers, who may work at different times, and contribute different amounts of work to the project itself (Grune, 1986). It then caught on because it's practical, and it scales well: so well that something designed for three people is now used to coordinate thousands of contributors. The following describes the rationale behind CVS:

One person may be working a steady ten hours a day on the project, a second may have barely time to dabble in the project enough to keep current, while a third participant may be sent off on an urgent temporary assignment just before finishing a modification. It would be nice if each participant could be abstracted from the vicissitudes of the lives of the others.

(Grune, 1986, p. 1)

Wikipedia's editing process is very similar: a central body of text (the article) is modified in parallel by any number of editors, who modify a local copy of the article, temporarily displayed in their browser, and then submit the change to the article. Every edit generates a new version of the article, which is then displayed, while older versions of the article are archived and kept visible in its history. History pages allow comparisons between any two versions of the same article, since its creation.

When it comes to knowledge-production on Wikipedia, the method of versioning achieves, in my view, two aims: it reduces perceived sources of bias and sets up a system of automatic correction of errors. As I will evidence more thoroughly in section 4.4, bias is a persistent worry for Wikipedians, who have, over time, put various systems in place in order to reduce it. Versioning addresses worries around bias by reducing the involvement of editors, in two ways: by forcing small contributions, and by minimising interaction between contributors. Forcing to update the page in small increments spreads the power thinly among the editor community: no one can make major changes directly, without going through some kind of consultation with the wider community. Larry Sanger explicitly links technical affordances and limiting the power of individual editors:

Wiki software also discourages (or at least does not facilitate) the exercise of authority, since work proceeds at will on any page, and on any large, active wiki

it would be too much work for any single overseer or limited group of overseers to keep up.

(Sanger, 2005).

Sanger frames piecemeal editing as an anti-authoritarian feature, in line with traditional hacker or free software ethics. Here, I focus on the epistemic fallout: the inability, for any given editor, to shape the angle of a whole page. If one wanted to simply relate facts in neutral language, a way to express one's views would be to juxtapose those pieces of information in such a manner that they create a narrative with a clear conclusion, or designed to raise certain questions, much like a documentary filmmaker would do. Piecemeal editing doesn't eliminate narrative – Wikipedia is a knowledge-production enterprise exactly because it builds a narrative; the connections between facts are what makes Wikipedia's content non-neutral – but places it in the hands of a collective, as opposed to individuals. As a result, it is much more difficult for any single point of view to be deliberately written into an article. In the context of Wikipedia's ethos, eliminating individual points of view is sufficient to clear the bar for neutrality – defined, it might be worth reminding here, as “neutral point of view”.

At the same time as reducing the potential imprint of individual editors, versioning allows coordination by establishing a relation between each editor and the text, making direct communication between editors redundant. Automatic coordination eliminates the need to agree on practicalities such as timings or priorities, wait for others to finish their work, or even think about other collaborators at all. In so doing, versioning reduces contributors to contributions. Versioning also foregrounds the relationship between the contributor and the product, to the detriment of the relations between contributors themselves. Minimising interactions, finally, prevents contributors from making any joint efforts to enact large-scale change.

The thrust towards reducing the influence of the human element in the editing process is, in and of itself, a form of validation, appealing to mechanical objectivity (Daston & Galison, 2010):

By mechanical objectivity we mean the insistent drive to repress the willful intervention of the artist-author, and to put in its stead a set of procedures that would, as it were, move nature to the page through a strict protocol, if not automatically. This meant sometimes using an actual machine, sometimes a person's mechanized action, such as tracing.

(Daston & Galison, 2010, p.121)

Daston and Galison are concerned with visual representations of nature, here, but the concept applies to Wikipedia as well, insofar as editors are aiming for as faithful as possible a representation of reality, through reproduction of outside sources. The two methods mentioned in the quote – using a machine, or a person who acts like a machine – fully apply to Wikipedia.

The way Wikipedia was set up in the first place was reliant on automation: Nupedia, Wikipedia's precursor, was an encyclopaedia written by experts. It was very inefficient, producing an unsatisfactory number of articles. Wikipedia was created to involve non-experts that could enrich their work by crowdsourcing information. The role envisaged for Wikipedia editors was, essentially, the same as that of rambot. Wikipedia was designed originally as a machine for the production of

facts, reliant on a form of automation.

I have anticipated earlier that Wikipedia's editing process is equipped both with a system that downplays individual contributions, and an error-correction system. I have argued that versioning provides the former, and will now turn to the latter. Errors are addressed by Wikipedia's process through negative feedback, in the form of radical collaboration: the editing interface allows (almost) anyone to edit (almost) anyone else's work. Mistakes can, then, be easily corrected through communal effort. Radical collaboration provides a reason to trust Wikipedia content: regardless of how many mistakes might be made on a page, there is always a large group of people who may spot them and correct them. "Eventualist" Wikipedians place enormous trust in this mechanism: they believe that, given enough time, the self-correcting nature of Wikipedia will eventually guarantee complete accuracy (MetaWiki: Eventualism). Similarly, the philosophy known as Darwikinism frames Wikipedia's knowledge production process as an evolutionary process (Meta: Darwikinism).

If that is the case, then, the key grounding of Wikipedia's reliability is brute force. Collectively improving a product by assessing each other's work is common in coding culture, and described in an influential essay by Eric Raymond, called *The cathedral and the bazaar* (Raymond, 1998). In this essay, Raymond contrasts two models for collective programming. In the "cathedral" model, according to Raymond, workers are coordinated by a central bureaucracy, everyone does their job according to orders, the product takes a long time to make, and is released in its finished form. In the "bazaar" model, emerged in the context of Linux – hence, free software – development, a group of people work on the same product in a decentralised manner. Unfinished versions of the product are released, so that any errors can be spotted by others. In Raymond's words, "with enough eyes, all bugs are shallow", where "bugs" indicates mistakes in the code that cause issues to the functioning of software. Raymond's motto, that has been adopted widely in software production, is "release early, release often". Wikipedia operates on a similar model: the versioning-based editing system essentially means that all Wikipedia articles are unfinished products, available for improvement. The unfinished nature of Wikipedia is acknowledged in its editing policy (Wikipedia: Editing policy), and signalled on its interface as well: various elements of Wikipedia articles encourage the reader to fix mistakes, such as the "citation needed" tags, as well as banners that mark an article as a stub (draft) or warnings that an article has issues – both a warning to the reader, and a way of prodding them into improving it. The self-correcting nature of Wikipedia is one of the cornerstones of its credibility, and, counterintuitively, leverages the lack of expertise of its editors: lowering the bar for participation by eliminating the need for credentials means, in turn, that a large number of people can contribute; the more people are able to look at a page, the more likely they are to spot errors.

The technologies described above act on the behaviour of editors, by allowing or encouraging certain behaviours, such as piecemeal editing, which, in turn, has a bearing on the way in which Wikipedia's content is produced. Embedded in the design of these systems are some of Wikipedia's core values, such as neutrality and radical collaboration. These values were already embedded in the systems, designed to facilitate the production of software, that inspired Wikipedia's knowledge-production process. Thus, Wikipedia inherited, through its coded infrastructure – and other elements of its culture, of course, but I am concerned with technics here – an ethos that developed in a completely different environment, for an entirely different purpose than creating an encyclopaedia. Wikipedia's system of validation is both constituted of coded objects and based, in

part, on principles that apply to the creation of code. In the next section, I will describe the coded object in the curation of Wikipedia's content.

4.2.3. Curation with and through coding

Having shown how coded artefacts participate in the process of creation and validation of knowledge, I will now describe how they contribute to the curation of Wikipedia's content. By curation, I mean the activity of connecting articles, structuring their content into a coherent description of reality. Geiger (2014) describes the curatorial role of bots in the following terms:

bots make it possible to achieve a certain level of uniformity in style and content – that is, when Wikipedians all agree that, for example, reference lists should be in alphabetical order or articles should use American instead of British spelling. Bots also serve key governance roles, independently enforcing discursive and epistemological norms, particularly for newcomers.

(Geiger, 2014, p. 345)

Bots achieve curation through aesthetic and epistemic means: they make Wikipedia look and feel uniform, and filter out unwelcome contributions. Bots are made by editors and deployed on Wikipedia to make cosmetic changes to the text, such as checking the spelling of commonly misspelled words and ensure consistency in the variety of English (UK or American) adopted on a given page (Geiger, 2014; Wikipedia: Bots). I described in section 4.2.3 how bots are used to carry out the so-called “wikification” of text: that is, turning text into something that conforms to Wikipedia's standards and rules.

The epistemic role of coded objects – including, but not limited to, bots – consists among other things in spotting vandalism, defined on Wikipedia as, “any addition, removal, or modification that is intentionally humorous, nonsensical, a hoax, offensive, libellous or degrading in any way” (Vandalism on Wikipedia). These tools report, to human editors, content that may infringe copyright and contributions by users that have been banned (Wikipedia: Types of bots).

These tools leverage various kinds of information, including the content of the edit, metadata, and human assessments. I will list a few here, for illustrative purposes. Huggle, released in 2008, uses machine learning techniques (based, among other things, on assessments made by human Wikipedia editors) and features a white list of trusted editors (Wikipedia: Huggle). It evaluates whether an edit appears to be vandalism, and allows editors to quickly revert suspicious edits (Wikipedia: Huggle). STiki, now inactive, used information about the edit to, ultimately, assess the reliability of the author; it calculated vandalism probabilities on the basis of information such as the author's status within the community and geographical origin of the edit (Wikipedia: STiki). Wikitrust, now inactive, used to measure the reliability of a page's content by examining its edit history (*ibid.*). I will provide more detail in section 4.3.3; it is important to note here that Wikitrust has been criticised for the way the calculation of reliability scores was done (WikiTrust). ORES (Objective Revision Evaluation Service), as described by its developers, “works by training models against edit- and article-quality assessments made by Wikipedians and generating automated scores for every single edit and article.” (Halfaker and Taraborelli, 2015).

It is important to note how the tools listed above intersect with Wikipedia's social structure: if

they are extensions of their designers, whose power do they amplify? The answer to this question is not straightforward. First of all, designers of anti-vandalism tools can be volunteers, Wikimedia Foundation Engineers, or a combination thereof: Huggle was developed by a Wikipedia editor (User: Gurch); Wikitrust at the Online Collaboration Lab at the University of Santa Cruz (WikiTrust); ORES was developed in-house, by the Wikimedia Foundation, but volunteers have been encouraged to contribute, with the explicit aim of encouraging diversity and openness, and in response to criticism leveraged to older tools (Halfaker & Taraborelli, 2015). To complicate matters, the usage of these tools varies depending on user: Huggle should only be used by editors with specific permissions (Wikipedia: Huggle); ORES allows users to pick a specific model to evaluate edits (Halfaker & Taraborelli, 2015).

Coded artefacts have a pervasive role in curating Wikipedia's content, and entertain complex relationships with human actors. By evaluating contributions based on markers of reliability defined by their designers, anti-vandalism tools operationalise their designers' epistemologies. They also absorb assumptions of human editors through training data. Finally, not everyone within the Wikipedia community is equally equipped to design or use these tools, but the lines of difference are intricate, trailing the uneven terrain of permissions, technical skill, status. The question of power and the use of technology is an important one, and one that spills beyond the detection of vandalism. In the next section, I will extend my reflection on these matters, laying the groundwork for further analysis of the role of coding practice in defining Wikipedia's epistemology.

4.2.4. With or through?

At this stage, it seems fitting to wonder about the relationship between Wikipedians and the technologies described above: are Wikipedians making the rules by coding technology? Or is technology constraining the activity of Wikipedians? Where does power reside, in this scenario? The creation and deployment of wide-ranging tools to curate the encyclopaedia raise the important question of how much of the "bazaar" spirit mentioned in section 4.2.2 is actually alive on Wikipedia. Encyclopaedias have, historically, been divided between planned and unplanned, where planned encyclopaedias started with a general structure, created by editors, that was then gradually populated with content, and unplanned encyclopaedias, updated gradually, in a less directed way. Which group does Wikipedia belong to?

Initially, Wikipedia as a whole was definitely more similar to a bazaar than a cathedral, in true decentralised software development spirit: Larry Sanger describes in his memoir how he set out a few, basic principles, and trusted editors to respect them. Over time, however, partially due to habit, and partially to controversies that pushed Wikipedia's then-management to introduce limitations to what could and could not be edited, structures have emerged, which both constrain the work of editors (e. g. protection of certain pages) and create specific spaces for content to be produced (e. g. categories). In some cases, such as protection, structures are prescriptive: if a page is protected, only certain categories of people can edit it. In other cases, they act more like guides: categories invite editors to organise content in a specific way, because it is easier, for new work, to fit pre-existing parameters than to create new ones. As I described in chapter 1, alongside epistemic structures, code has solidified over time as well. A cathedral was then, so to speak, cobbled together from the output

of a bazaar¹¹.

Eventually, Wikipedia came to assume the appearance of a work of outsider architecture: built without planning, workable, but somewhat haphazardly structured and impossible to change radically without fear of accidentally tearing it down. Around this central building, the bazaar is, however, still active: adjustments and extensions are being built that, while not changing the essence of the edifice, try to make up for its shortcomings, or marginally change its layout.

Wikipedia, then, doesn't fit neatly in the planned or unplanned category. It is somewhat reluctantly planned, and in a way that doesn't involve a designed structure gradually filled with new content. Rather, the way Wikipedia's coded infrastructure influences content can be characterised in terms of protocol, as defined by Alexander Galloway (2004):

protocol is more than simply a synonym for "the rules." Instead, protocol is like the trace of footprints left in snow, or a mountain trail whose route becomes fixed only after years of constant wear. One is always free to pick a different route. But protocol makes one instantly aware of the best route—and why wouldn't one want to follow it?

Thus, a better synonym for protocol might be "the practical," or even "the sensible." It is a physical logic that delivers two things in parallel: the solution to a problem, plus the background rationale for why that solution has been selected as the best. Like liberalism, or democracy, or capitalism, protocol is a successful technology precisely because its participants are evangelists, not servants. Like liberalism, democracy, or capitalism, protocol creates a community of actors who perpetuate the system of organization. And they perpetuate it even when they are in direct conflict with it.

(p. 245)

Protocol, according to Galloway, is a modulating force, tracing paths for information and people through its very existence. It is a fitting description for Wikipedia's infrastructure: a decentralised lattice of technologies and conventions shape Wikipedia's content and activity, in a way that is not entirely prescriptive, nor entirely free. Wikipedia's infrastructure, however, is more than a top-down imposition of rules embodied into code. It is also deeply intertwined with Wikipedia's ethos and community practices.

What distinguishes Galloway's account of material agency from other accounts, is that protocol doesn't just embody a rule – it also acts in such a way that changes the subjectivities of those around itself, inducing a desire to follow it. Protocol legitimises itself by maintaining, with the community around it, a relationship based on desire.

So far, within the literature, the implicit assumption has been that Wikipedians carry out their epistemic activities on top of and around a bedrock of coded infrastructure, made by someone other

¹¹ Bazaars, of course, don't produce anything, and using a bazaar as a metaphor for production is only one of the many issues with Raymond's essay, which would be uncharitable to list here. Nonetheless, I am keeping this nomenclature here for consistency, and to show the connection between programming lore and Wikipedia's culture, while acknowledging it may not be particularly illuminating per se.

than editors, which is used as a springboard to create knowledge. Sometimes Wikipedians use infrastructure as it was intended; other times, intentionally misuse it in order to game the system (Ford & Geiger, 2012; Ford, 2022). When bots are mentioned, they are framed in terms of extension of power – as something that allows a specific human actor to do more (see e.g. Geiger, 2014).

Here, I want to bring attention to two facts: that Wikipedians, as a collective, are coders themselves; and that the boundaries between volunteers and the Wikimedia Foundation staff are porous and shifting. The fact that Wikipedians are coders themselves changes their relation to coded objects: the two main options described above – use-as-intended and purposeful misuse – don't cover the full range of possibilities. A third option needs to be added: intervention, through creation or modification. As a consequence, and combined with other features of Wikipedia's infrastructure defined above, epistemic practice accrues a significant technical component. Wikipedians create knowledge through code. Take bots: epistemic advances can be achieved through the use of bots, but bots are also designed based on epistemic values and assumptions.

However, one might object, some of the coded objects through which Wikipedians produce knowledge – the editing interface, ORES – were not created by volunteers. What does this mean, in terms of agency? The first thing I'd say is that Wikipedians are also involved in updating objects like ORES. As mentioned in section 4.2, the Wikimedia Foundation engineers routinely enlist the help of volunteer programmers to work on the platform. Secondly, even those who don't have engineering abilities can put pressure on the engineering team by writing on their talk pages. The practice of communicating directly with members of the Foundation is bolstered by the culture of openness and transparency inherited from the free software movement. As the Principal Software Engineer put it during our interview:

It's also it's an interesting dynamic to you know because... I'm not saying that to say that we are caving to yells or whatever, but I don't know how many other organizations are out there where the CEO, the board, are all on wiki in a way, where random people can come on their talk page and yell at them.

(Moriel Schottlender)

The landscape becomes more complex depending on the notion of authorship we decide to adopt. If we consider authorship to be individual, then there is a considerable difference between someone adding a citation through the editing interface and David Ramsey, who coded a bot to write articles for him. However, if we consider authorship to be collective – which, in Wikipedia's case, seems justified, given the ethos of radical collaboration and anonymity that animates the project – then things change: it doesn't matter who codes what. The community, composed of all volunteer editors, produces, validates, and systematises knowledge through coding. The various technologies that constitute the bones of Wikipedia's production system are embedded in practice, both because they are, in part, made by Wikipedians themselves to regulate their own work, based on solving problems as they emerge in everyday activities, and because they embody a shared understanding of knowledge, that Wikipedians, broadly speaking, and with some caveats which will be explored in section 4.5, want to uphold.

4.3. Sorting technologies

In *Writing the Revolution* (2022), Ford analyses the Article Wizard, a tool developed to allow unregistered users to create new articles, to be reviewed before publication by registered users. While the Wizard is designed to simplify article creation, it also acts as a filter for contributions of unregistered users. First, the Wizard constantly reminds users of the basic principles of Wikipedia editing, and, if these are not respected, re-routes the editor towards other ways of contributing to the encyclopaedia (Ford, 2022). Secondly, the Wizard makes it impossible for unregistered editors to contribute articles that relate to recent events (*ibid.*), probably due to the sensitive nature of covering issues that are in the news. Third, using the Article Wizard is in and of itself a way of attracting scrutiny, as registered users are prompted to evaluate articles created in this way, and are aware of the low level of experience of the user who made the article (*ibid.*). In sum, despite intentions, the Article Wizard may make it more, not less, difficult for unregistered users to contribute.

Ford's analysis highlights various sorting mechanisms deployed by the Article Wizard: reminding unregistered users of their ignorance of Wikipedia's practices; preventing them from contributing to articles with ample readership, and holding their contributions to arguably higher standards than those of more experienced editors. A few different mechanisms are at play here, leaving smaller and smaller latitude for counteraction from editors: dissuasion can be overridden; surveillance may increase likelihood of, but doesn't guarantee, the impossibility to contribute; while impeding access is an insurmountable barrier to editing.

In this section, I will carve some detail into Ford's analysis, by disentangling the various mechanisms through which coded artefacts operate distinctions between suitable and unsuitable contributions, and good and bad contributors. Based on a set of case studies, I will describe four mechanisms: explicit sorting, operated by artefacts made to discriminate between contributions or contributors; sorting through affordance, concerning artefacts that were not designed to discriminate, but do so nonetheless because of how they function; sorting through material constraints, occurring as a result of the material qualities of specific objects; and sorting by tending to the conservation of Wikipedia's culture.

4.3.1. Explicit sorting

Some coded objects used on Wikipedia operate an explicit sorting function on contributions, in a deliberate, rule-based manner. That is, they are designed to discriminate between two contributions, determine which is correct, and do something about it – either delete the contribution (revert the edit) or alert a human editor.

I have already introduced these tools in section 4.2.3, as curation tools. Here I wish to unpack their role as actors within Wikipedia's ecosystem, by looking at the criteria they use to discriminate between good and bad contributions. The analysis laid out here will allow me to describe, in section 4.4.1, the concept of expertise embedded in Wikipedia's coded infrastructure.

Anti-vandalism tools, as described above, base their evaluations on metadata related to the edit, the content of edit itself, or a combination of the two. Some gauge good faith by using machine learning techniques. Tools based on metadata tend to use it as a proxy for markers of identity, ultimately evaluating the likely trustworthiness of the editor. WikiTrust, now discontinued, used to

look, for each word that was part of an edit, at the place where the word was, the author (who included the word in the text), and the level of “trust” associated with the word (WikiTrust). In turn, the “trust” metric relied on the judgment of editors that were considered to be “high authority” (*ibid.*). If an edit was left on the page for a long time, it was considered more trustworthy (*ibid.*). STiki “examines only four fields of an edit when scoring: (1) timestamp, (2) editor, (3) article, and (4) revision comment. These fields are used to calculate features pertaining to the editor’s registration status, edit time-of-day, edit day-of-week, geographical origin, page history, category memberships, revision comment length, *etc*” (Wikipedia: STiki). Huggle employs, among other techniques, a whitelist of reliable editors (Wikipedia: Huggle). To a large extent, then, these tools have the explicit function of endorsing edits made by experienced Wikipedians. This is done intentionally, as explained by the creators of STiki:

Intuitively, one would expect long time members of the Wikipedia community to be vested in its growth and familiar with its policies. Conversely, malicious editors may employ a Sybil attack, creating temporary accounts to abuse the associated benefits. Time-since registration should encode such behaviors.

(*West et al.*, p. 3)

The “intuitively” at the beginning of the quote above does a lot of rhetorical heavy-lifting, and uncovers the mechanism by which epistemic assumptions held by developers are operationalised into technology: behaviours attributed to experienced editors are used as a basis for employing a specific technique in assessing the likelihood of disruption. In turn, basing the evaluation on time from registration makes it more likely to spot vandalism perpetrated by newer editors, thereby confirming initial assumptions and feeding the notion that new editors are disruptive.

Another epistemically charged way of detecting vandalism is the use of spatio-temporal data. STiki uses timestamps and geolocation (through IP addresses) in order to determine the likelihood for a given edit to constitute vandalism, effectively putting a geographical bias, enforced by a combination of spatial coordinates and time zones, on the survival rate of new edits. The decision to use space and time as variables is a combination of its designers’ intuitive understanding of how vandalism works, and STiki’s design history: the tool was modelled on email spam filters (West *et al.*). The potential implications for the diversity of Wikipedia’s editor base are clear. First, as a baseline, American editors are trusted above everyone else. Subtler effects might also occur as a result of work patterns. STiki posits that more unhelpful edits are done during the week than at the weekend. There might be implications, here, regarding the class of the editors, and their professional status. It does tally with the affinity between Wikipedia’s ethos and white-collar work. The way in which Civility is phrased, for instance, implies an audience that is familiar with the norms of office work:

Be professional. Just because we are online and unpaid does not mean we can behave badly to each other. People working together in a newspaper office are

not supposed to get into punch-ups in the newsroom because they disagree about how something is worded or whose turn it is to make the coffee.

(*Wikipedia: Civility*)

Article talk pages should be, on the whole, considered to be professional workspaces. They're places to talk about how to improve the article, and to discuss the article (though it's OK for conversations to wander into related areas, or go more into depth than the article does, as that helps with research and gives ideas on improvement). But an editor's talk page is more like their kitchen; it's more informal, and (within reason) it's up to them what happens in there.

(*Wikipedia: Civility*)

As I argued in chapter 3, Wikipedia's aesthetic also connects trustworthiness with white-collar professionalism: increasing the likelihood of flagging as vandalism contributions performed by shift or blue-collar workers plays into the same worldview.

As well as metadata, anti-vandalism tools can evaluate the content of edits. It is very difficult for an automatic detection tool to assess whether an edit is malicious *per se*, with the exception of blatant breaches, such as the use of offensive language. Tools attempting to do so employ machine learning, trained on judgments of seasoned Wikipedians (e.g. ORES, Huggle). This makes sense from a practical point of view: long-standing editors know policies well. However, in an environment where policies are not always straightforwardly applied, and rather used as part of complex strategies (see section 1.3.2), the judgment of veteran Wikipedians may not be impartial. Machine learning-based assessment of vandalism, then, rewards gaming the system, and reproduces the preferences of those who already know how to.

Anti-vandalism tools then can, and do, favour edits by longstanding members of the community. At the same time, they indirectly sanction the work of new editors through another mechanism: paying more attention to categories and articles that tend to be controversial (as, for instance, STiki did). Controversial articles are often controversial because they cover politics, social justice, and other sensitive topics. Wikipedia covers world events as they happen. It has been proven that, when something consequential happens in the news, many inexperienced people flock to Wikipedia to support the editing effort (Zhang *et al.*, 2019). This is true, for instance, of the protests around Black Lives Matters and the Ebola outbreak (Li & Farzan, 2018). Consequently, and consistently with Ford's findings on the Article Wizards, edits of newcomers working on important, live issues are placed immediately on the firing line.

In sum, technical features performing sorting functions operationalise trustworthiness in ethically charged ways: collectively, they treat experienced editors as inherently more trustworthy, while at the same time displaying and, possibly, perpetuating an attitude of diffidence towards new editors and editors that don't fit the demographic profile of the majority of the Wikipedia community.

4.3.2. Sorting through affordance: litmus technologies

While tools operating explicit sorting are made to discriminate between people and/or

contributions, other coded objects, designed with different objectives in mind, do so by virtue of the affordances they offer their users. In this section, I will examine in these terms four Wikipedia-specific tools: the Visual Editor, the Source Editor, the editing process as a whole, and talk pages.

Affordances, as defined by Gibson (1979), are ways in which the environment lends itself to be used by specific animals: certain plants afford eating to certain animals, for instance. Affordances are a fundamental way in which the agency of objects is conceptualised in philosophy of technology, and particularly in postphenomenology (for a primer on postphenomenology, see chapter 2). Affordances, in the context of technology use, are one of the ways in which artefacts act on their environment – by lending themselves to certain uses. An important point about the notion of affordance is that it is a relational concept: affordances are not universal. For instance, a plant can count as food for a giraffe, but not for a lion. This variability is captured, as applied to technology in the concept of multistability, theorised by Don Ihde (1990). Multistability entails that the same artefact can afford different uses to different people, depending on their identity and background. For instance, a section of wikitext will afford editing only to those who know what wikitext is, what it does, and how to edit it.

I have introduced in section 4.2.1 the two different interfaces that allow to edit Wikipedia articles: the Source Editor and the Visual Editor. I argue here that, as a set, they fulfil a sorting role by exploiting multistability: they afford different behaviours to different people. The multistability of the two editors is built on epistemic, aesthetic, and affective bases, which, collectively, grant a higher level of epistemic credibility to tech-savvy editors.

First of all, their appearance is significant: the Source Editor displays wikitext (Fig. 4.1), while the Visual Editor shows rendered text, similar to the article page (Fig. 4.2). While wikitext is technically new to anyone who edits Wikipedia for the first time, its appearance and logic will be familiar to those who know HTML. It was probably quite easy to understand for the pioneer editors of Wikipedia, who were, after all, the same people who programmed it. Additionally, HTML editing was a relatively widespread skill in the early 00s: those who participated to online communities were at least acquainted with markup, because it was commonly encountered in the use of popular

Editing Walnut

✱A Add languages ▾

Article Talk

[Read](#) [Edit source](#) [View history](#) [Tools](#) ▼

 You are not logged in. Your IP address will be publicly visible if you make any edits. If you [log in](#) or [create an account](#), your edits will be attributed to a username, among [other benefits](#).

 Content that **violates any copyrights** will be deleted. Encyclopedic content must be verifiable through **citations to reliable sources**.

```
B [ ] [ ] [ ] [ ] [ ] [ ] > Advanced > Special characters > Help > Cite
{{Short description|Nut of any tree of the genus Juglans}}
{{about|walnuts as food|the genus|Juglans|walnut wood (timber)|Juglans#Wood|other uses}}
{{Use dmy dates|date=January 2023}}
[[File:Walnuts - whole and open with halved kernel.jpg|thumb|Walnuts]]

[[File:Inside of a walnut in growth.jpg|thumb|upright|Inside of a walnut in growth]]
[[File:Three-part Walnut-8283.jpg|thumb|upright|Three-segment walnut]]
[[File:Juglans regia Echte Walnußfrucht 2.JPG|right|thumb|Walnut shell inside its green husk]]
[[File:Adriaen Coorte - Two Walnuts.jpg|thumb|right|Artistic depiction of two walnuts ([[Adriaen Coorte]], 1702)]]
A '''walnut''' is the edible [[seed]] of any tree of the genus '''[[Juglans]]''' (family [[Juglandaceae]]), particularly the Persian or English walnut, '''[[Juglans regia]]'''. They are [[accessory fruit]] because the outer covering of the fruit is technically an [[Involucral bract|involucre]] and thus not morphologically part of the carpel; this means it cannot be a drupe but is instead a drupe-like nut.

After full ripening, the [[husk|shell]] is discarded, and the kernel is eaten. Nuts of the eastern black walnut ('''[[Juglans nigra]]''') and butternuts ('''[[Juglans cinerea]]''') are less commonly consumed.

== Description ==
Walnuts are the round, single-seed [[Drupe|stone fruits]] of the walnut tree. They ripen between September and November in the northern hemisphere. The brown, wrinkly walnut shell is enclosed in a [[husk]].<ref name="GKH">{{Cite web |last=Grant |first=Amy |date=19 April 2021 |title=Walnut Tree Harvesting: When Are Walnuts Ready To Pick |url=https://www.gardeningknowhow.com/editable/nut-trees/walnut-walnut-tree-harvesting.htm |access-date=4 December 2021 |website=Gardening Know How}}</ref> Shells of walnuts available in commerce usually have two segments (but three or four-segment shells can also form). During the ripening process, the husk becomes brittle and the shell hard. The shell encloses the kernel or meat, which is usually in two halves separated by a membranous partition.<ref name="GKH"/> The seed kernels – commonly available as shelled walnuts – are enclosed in a brown [[seed coat]] which contains antioxidants. The antioxidants protect the oil-rich seed from atmospheric oxygen, preventing [[rancidification|rancidity]]<ref name="Wd1" />
```

Figure 4.1. Source Editor – “Walnut” edit page.

Walnut

🌐 Add languages ▾

Article Talk

[Read](#) [Edit](#) [View history](#) [Tools](#) ▼

↶ ↷ Paragraph ▾ A ▾ 🔗 🗣️ Cite ☰ ▾ Insert ▾ Ω

[Publish changes...](#)

Short description

+ Insert paragraph

This article is about walnuts as food. For the genus, see [Juglans](#). For walnut wood (timber), see [Juglans § Wood](#). For other uses, see [Walnut \(disambiguation\)](#).

 Use dmy dates

A **walnut** is the edible **seed** of any tree of the genus *Juglans* (family *Juglandaceae*), particularly the Persian or English walnut, *Juglans regia*. They are **accessory fruit** because the outer covering of the fruit is technically an **involucre** and thus not morphologically part of the carpel; this means it cannot be a drupe but is instead a drupe-like nut.

After full ripening, the **shell** is discarded, and the kernel is eaten. Nuts of the eastern black walnut (*Juglans nigra*) and butternuts (*Juglans cinerea*) are less commonly consumed.

Description

Walnuts are the round, single-seed [stone fruits](#) of the walnut tree. They ripen between September and November in the northern hemisphere. The brown, wrinkly walnut shell is enclosed in a [husk](#).^[1] Shells of walnuts available in commerce usually have two segments (but three or four-segment shells can also form). During the ripening process, the husk becomes brittle and the shell hard. The shell encloses the kernel or meat, which is usually in two halves separated by a membranous partition.^[1] The seed kernels – commonly available as shelled walnuts – are enclosed in a brown [seed coat](#) which contains antioxidants. The antioxidants protect the oil-rich seed from atmospheric oxygen, preventing [rancidity](#)^[2]

Walnut trees are late to grow **leaves**, typically not doing so until more than halfway through the spring.

Chemistry

Walnut hulls contain diverse [phytochemicals](#), such as [polyphenols](#), that stain hands and can cause skin irritation. Seven phenolic compounds, including [ferulic acid](#), [vanillic acid](#), [coumaric acid](#), [syringic acid](#), [myricetin](#), and [juglone](#), were identified in walnut husks; juglone had concentrations of 2–4% fresh weight.^[3]




Figure 4.2. Visual Editor – “Walnut” edit page.

platforms such as WordPress and MySpace, and even the first online forums. Over time, HTML editing became a more specialised skill (Kennedy, 2012), while, contextually, interfaces became easier to use, and custom designs fell out of fashion, replaced by the uniform aesthetic of platforms such as Facebook and Twitter. In sum, the Source Editor is more user-friendly for those who are familiar with HTML.

There is also a subtler point to be made about wikitext, that has to do with cultural meanings associated with code and the practice of coding. As anticipated in section 4.2.1, whether markup constitutes a programming language is a contested issue. The reasons why it's contested are related

to what coding represents, and the cultural capital ascribed to those who can code. It is worth briefly explaining the dispute here, which centres on HTML, but is applicable, by extension, to other markup languages.

Sources disagree on whether HTML is a programming language. Everyone agrees, however, on the fact that HTML, which stands for Hypertext Markup Language, is a markup language: that is, a system of symbols – letters, numbers, and punctuation – that allows to tag a text document, and define how each part of the text will be displayed in a Web browser (Hemmendinger, 2023). Some elements of the document will be displayed as text, others as images, videos, graphic elements and so forth. The question is: is telling a Web browser what to do with text sufficient to qualify as a programming language? This is difficult to establish because the definition of what a programming language is and what it does varies; differences in definition often vary due to cultural, rather than technical, factors: definitions appear to be created precisely to include or exclude, through technical specifications, certain languages that are culturally deemed inferior.

An indication that the reason why definitions vary is often cultural rather than technical is that the definitions themselves display few continuities, with the exception of pointedly excluding or including HTML. In other words, the definitions written by those who are happy with calling HTML a programming language all look the same; while those who want to exclude HTML from the elite of programming languages call, each, upon a different reason to do so. Some will, then, say that a programming language is something that interacts with the CPU of a computer, and as such it must contain variables, if statements, and loops (e.g. DeClute, 2022). Others will say that a programming language needs to be Turing complete, meaning it needs to be able to run any program that Alan Turing’s machine used to be able to run (e.g. Abramowski, 2022).

These debates are largely informal, and occur among technology enthusiasts, rather than academics: whether HTML is formally considered a programming language is largely irrelevant in practice. What matters is how to use it. Why do some feel so strongly about it then? It has been argued that the reason is that HTML has historically been easy to use, and used by amateurs; and in an environment, such as coding communities, where prowess defines one’s status and identity, excluding HTML is a matter of defending one’s identity, or, in other words, gatekeeping (for a summary of this perspective, see Stieg, 2022).

I am not going to attempt to settle the dispute described above. I have already elaborated on the material affinity between markup and code (section 4.2.1). Here I want to draw attention to the cultural affinities between markup and code, and how they interact with its materiality. If code is text that does things, programming is framed as the powerful practice that creates text that does things, and markup is somewhat associated with programming, then those who are new to markup may feel like they need a certain level of skill in order to intervene. There is also an aesthetic similarity: wikitext looks a lot like code, especially to novices. Of course, any programmer will easily see that the structure of wikitext is very different from, say, Python: it doesn’t contain variables, or if statements, and it contains tags.

But someone who has never coded will see a page filled with text that does things, in a way that they don’t understand. Coding literacy is not just about skill, it’s about being comfortable with digital technologies: those who have not grown up tinkering or playing videogame may be scared of coding (Vee, 2017). This is particularly true in a world where programming is often regarded as a dark art,

and programmers enjoy great social prestige and reverence, the uninitiated might find the mere appearance of a code-like page to imply that they don't have the skills, or even the intelligence – or any such analogous construct that implies innate ability – to edit wikitext. The common novice programmer fear of breaking something, especially when intervening on pre-existing code, presents itself with markup as well. This fear can be off-putting to newcomers, especially in a culture, like Wikipedia's, where errors can elicit blunt or aggressive responses.

Any anxieties around editing markup may be compounded by the fact that Wikipedia's interface presents frequent reminders of the consequential nature of editing, especially related to articles that are likely to be seen by many (see Fig. 4.3 for an example). The article in Fig. 4.3 (below) doesn't contain a warning because of its sensitive content, or fundamental epistemic importance, but because, on the day I opened the edit page, it was featured on Wikipedia's Main Page, hence likely to be widely read: the warning reifies the social fear of failing in front of a large audience. The newcomer's anxiety is mirrored by the anxiety of those who designed the warning – the fear that novices might do damage.

Editing with the Source Editor is not coding, but provides, in practice, a coding-adjacent experience, with all that entails. It has ethical and affective consequences, that reverberate in the whole of the knowledge-production process. Commonly cited issues, within the Wikipedia community, are the slow uptake of new editors, the way new editors are mistreated by more experienced Wikipedians, the difficulty of operating change, and epistemic injustice. I am not saying that the code-like nature of the Source Editor is the sole responsible for any of these complex problems, but it does have an impact, by contributing to an atmosphere of anxiety around editing. The Visual Editor and the Source Editor are not just visually different: they differ in terms of the actions they allow. The Visual Editor's functionality is not as advanced as the Source Editor's: editing

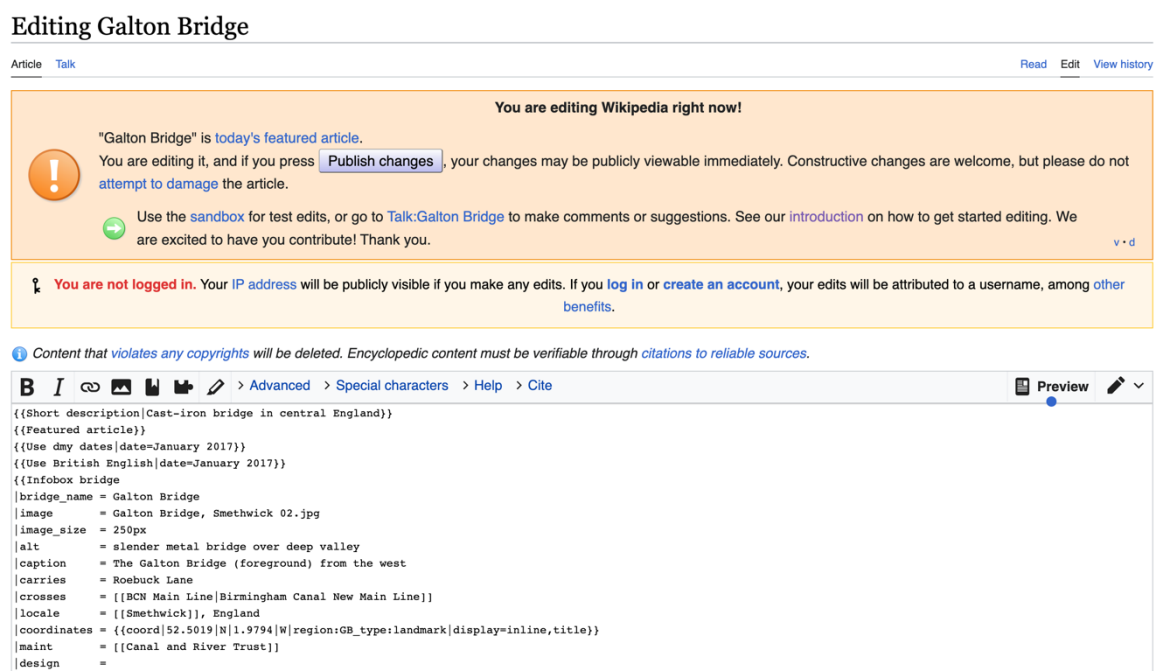


Figure 4.3. Galton Bridge edit tab, on the day it was a Featured Article.

tables and small sections is more difficult, some actions are impossible, loading times are slower, and

more (Wikipedia: VisualEditor).

The two editors are also framed differently on the website, inviting Source Editors to act, and Visual Editors to caution. Documentation about using the two systems differs in its assumed audience. In the source editor introduction summary, there is a call to action (“Put what you learned into practice”, followed by number of articles that need citations), which is much more cautious in the Visual Editor guidelines (“Test what you’ve learned in a sandbox”, i.e. a test space where no changes can be made to Wikipedia’s text). Source editor instructions contain relatively fine-grained guidance about editing, while Visual Editors instructions provide only a general overview – thereby assuming the user to be less knowledgeable about the editing process – and entry-level guidance (“Your first Wikipedia article”). In the section about images, source editor users are invited, over and over again, to upload images, while Visual Editor users are reminded of the wealth of already-available images and taught how to use them. A call to action is provided for Source Editor users, but not Visual Editor users. Source Editor users are invited to register, Visual Editor users are not. In a very concrete way, Source Editor users are assumed to know more about Wikipedia, and, crucially, about a host of other matters, such as choosing and uploading images; they are also deemed to be a good fit for the community. Each editor carries assumptions around technical prowess, the connection between technical prowess and other skills, and between skill or knowledge and potential standing in the community.

The hinge of the reciprocal connections between technical skill, knowledge, and belonging is, in my view, the figure of the geek, or nerd (I am using the terms interchangeably here). A geek, in modern terms, is someone who typically combines tech savvy, or even coding skills – often self-taught, whether applied in an amateur or professional context – and a high level of intelligence, and/or wide general knowledge (Kendall, 1999a). Nerds are characterised by curiosity, and the ability to find information about a topic they are deeply interested about – a key talent of Wikipedia editors. Implicitly conceptualised as geeks, Source Editor users are trusted more than Visual Editor users, both in terms of technical expertise and likelihood of contributing meaningfully to the encyclopaedia. Nerds are also characterised by specific traits and interests shared with others, which mark them as members of a specific subculture. Technical expertise is, then, assumed as a proxy for encyclopaedia-editing skills, and both are markers of potential cultural fit.

While the two editors discriminate on the basis of technical savvy, the editing system privileges those who are already familiar with Wikipedia’s unspoken rules. As I pointed out in section 4.2, because of the way the editing system is designed, only small, piecemeal edits are allowed. That means that, by design, no one has the power to make large changes – which ties in with NPOV: if everyone makes small contributions, the article is not from any one point of view. The curation of the content of a Wikipedia article is delegated, to a large extent, to the process.

At least, that’s true in principle. In practice, those who know how to game the system have ways of imposing their point of view on an article. An example of how this can be done is by using citation needed tags. As explained in chapter 3, a citation needed tag is a superscript, placed after a statement, which reads “(citation needed)”. It indicates that the statement is not currently supported by a reference, and is an implicit invitation for editors to find a reference for that statement. It is also an implicit, if low-level, criticism: Wikipedians are supposed to support edits with citations, and highlighting the absence of a citation is equivalent to noting that a fellow editor hasn’t done a good

enough job. In and of itself, one “citation needed” tag is not a major attack – just a light reproach. However, when many “citation needed” tags are added in a row – especially when they are added by the same editor – that means the editor is trying to signal that an extensive intervention is needed. While littering a page with citation needed tags is contrary to the spirit of NPOV, it is also technically allowed: there isn’t a built-in limit to how many “citation needed” tags can be added by one person in the same article – while, for instance there is a cap on reverts. Allowing the placement of multiple citation needed tags lets those who are familiar with the subtext of using citation tags to game the system. Experienced Wikipedians, once again, are able to exert influence by virtue of their belonging to the community.

Talk pages are also prime gaming the system ground. In chapter 1, I have highlighted how discussion on talk pages is highly regulated. As I mentioned then, the relationship between regulation and behaviour is not straightforward: it’s not simply a case of applying rules and sanctioning those who don’t follow them. There are some caveats. Explicitly citing policies (Wikipedia: Wikilawyering) is frowned upon, and one can always ignore rules if they don’t serve the mission of the project (Wikipedia: Ignore All Rules). The rule of thumb, within the community, is that anything is allowed, unless and until someone objects. Arguments on talk pages are mostly won by default – when objections stop coming. Behaviour, much like editing, is a negative-feedback system. As a result, those who are not familiar with local culture struggle to make meaningful contributions, because they cannot foresee, or in some cases even make sense of, negative reactions from more established editors (Ford & Geiger, 2012).

The way Wikipedians use coded infrastructure to game the system challenges traditional accounts of technology, that tend to separate use and misuse, where misuse is sometimes a repurposing of a certain object, in a way that was either unforeseen, or even opposed to the use intended by the designer. In Wikipedia’s case, is applying multiple citation needed tags a correct use of the technology, or a misuse? I suggest it’s neither. It’s an action that transcends use and misuse, because of the nature of prescriptiveness within Wikipedia’s community. If Wikipedia were a strict, rule-based community, then it would make sense to speak of use and misuse: citation needed tags would be meant exclusively as means of flagging missing information, and those who use them otherwise would be breaking a rule. But Wikipedia is not regulated in a traditional sense – as articulated above, Wikipedia’s rules are made to be ignored by those who know how to take advantage of ignoring them. Creative use of elements of the interface follows the same logic: it is a use that is neither allowed nor forbidden – in the spirit of do-ocracy, if you do it, and get away with it, then it’s ok to do it. And getting away with it is easier for those who belong to the group of people who know the rules of the game, and when and how to break them. Whether something is used appropriately or not doesn’t depend on the action itself: it depends on the social context and fallout of that action.

Talk pages and the editing interface, then, as well as allowing for certain actions (talking and editing) to be performed, act as a sorting mechanism, dividing editors and potential editors into two categories: those who are well-equipped to contribute to the encyclopaedia and those who are not, both in terms of skills and belonging to the community. Talk pages and the editing interface highlight the compatibility between one’s personal features and those of the ideal Wikipedian contributor, while allowing action only to those who pass the test. They are like puzzle boxes, which, in order to

be opened, need to be manipulated in specific ways. I am not saying that this happens intentionally – in fact, I don't have any reason to believe it does – but it does happen as a result of the Wikipedia's design, and its affordances. These specific functions of talk pages and the editing interfaces are deeply embedded in the process of editing Wikipedia.

4.3.3. Material sorting: wikitext and angle brackets

In some cases, sorting can occur far from the activity of human editors, through interactions between elements of the interface: rather than filtering people or contributions, features of Wikipedia's platforms interact by virtue of being close to each other, in ways that are epistemically significant. I will illustrate this point through an example that was explained to me by Moriel Schottlender, Wikimedia's Principal Software Engineer, in our conversation. It is a problem that, at the time of our conversation, was in the process of being solved; it is, nonetheless, illustrative of a consequential phenomenon that is likely to occur elsewhere as well.

The issue in question concerns the interaction between the functioning of the Visual Editor, the use of pages known as templates, and the material characteristics of markup. The Visual Editor works by overlaying an interface over wikitext: when an edit is performed through Visual Editor, behind the scenes, changes are translated into wikitext. In order to achieve this translation, the elements of an article page need to be structured in a way that Visual Editor can recognise automatically. To facilitate this process, templates are used. Wikipedia templates are "pages that are embedded (transcluded) into other pages to allow for the repetition of information" (Wikipedia: Templates). They are essentially blueprints for elements that are routinely used on pages: for instance, infoboxes – the frames displayed on the right side of many Wikipedia articles, displaying data about the article in a short, list-like form – are based on templates. The engineering team loaded all the Wikipedia templates into the Visual Editor to facilitate translation between interface and wikitext. Each template has a template data file attached to it, establishing what parameters the template accepts. This way, the Visual Editor user can simply use the visual interface to edit templates, through a dialogue box.

References, however, presented an issue. On Wikipedia, references are placed between ref tags, that look like this: `"Sunland<ref name="Cultivar" />"` (this is an example from the Walnut page). The Visual Editor allows the insertion of references through a dialogue box, and offers an option to repeat references as needed, if, as often happens, a reference is used multiple times on the same page. That means that the Visual Editor has to identify the things in the ref tags as references. The issue comes about with languages that are read right to left, such as Hebrew, Arabic and Farsi – in these languages, angle brackets are displayed the other way around compared to English, so `"Sunland<ref name="Cultivar" />"` will look like `"Sunland>ref name="Cultivar" /<"`, which in turn won't display a reference on the article page. In order to address this issue, a bot was created that changed reference tags to a template in the relevant language (Hebrew, Arabic, etc), which fixes the problem – but only to an extent. If someone wants to come back, later on, to the same page and add another reference through the Visual Editor, they won't be able to do it anymore, because the bot will have turned references into templates – they would not be recognisable as references anymore.

Unpacking this small case highlights a few ways in which material interactions can occur. First of all, the interaction is made possible by the fact that, as remarked above, code is text. In this specific

context, code is physically contiguous with inert text, which, as a consequence, means that acting on one will have repercussions on the other – hence whatever happens to inert symbols, happens to active text. In this case, active brackets are modified because inert text is modified. Secondly, the issue arises because programming and markup languages are modelled on the English language: had they been modelled on Farsi, of course, the situation would have been different. Third, the context in which this interaction occurs is one in which bots are active, which don't discriminate between the two kinds of text, perhaps because they are assumed, by native English speakers, to apply universally. Sorting, in this case, is the result of a complex system of interactions, that have occurred accidentally and unintentionally. Those who coded the Visual Editor, the bots, and Wikipedia's interface all had different aims, and assumptions about the world, that ended up penalising certain knowers.

4.3.4. Coded custodians

A recurring trait of the coded artefacts examined so far is their conservatism: they have a tendency to privilege what already exists, over deep change. Wikipedia's conservatism is not news, and unlikely to change, as it rests on the processes by which consensus is achieved (Reagle, 2010). Here I examine how said conservatism is embedded into technologies, that act as coded custodians: they keep the internal culture of Wikipedia stable, by amplifying the views of long-standing editors, discouraging newcomers, and re-affirming ideology.

The views of long-standing editors are amplified in various ways. First of all, bots carry with them their operators' values and preferences. Because of the system of approval for bots, long-standing editors are more likely to be able to deploy one. And because of the computing power that bots have and people don't, the values and preferences designed into bots inform large amounts of activity. Secondly, as I have detailed in section 4.3.1, the influence of experienced editors is further extended by using their assessments to train models that inform the activity of machine learning-based anti-vandalism tools.

The complementary tendency (and bolstering force) of the amplification of the views of long-standing editors is the proclivity to being harsh with newcomers. As detailed above, anti-vandalism tools tend to score the edits of newcomers as more likely to constitute vandalism, and do so explicitly and deliberately. As well as intrinsically protecting the status quo by preventing new people from joining the community by contributing to the project, these tools uphold a culture of prejudice against newcomers, which is, in itself, a known characteristic of Wikipedia's culture. It is so ingrained, that Wikipedia contains a behavioural guideline called Wikipedia: Please do not bite the newcomers (Wikipedia: Please do not bite the newcomers), reminding established editors that they used to be newcomers once, and inviting them to treat new additions kindly.

Finally, by filtering the potential contributor pool, these features preserve the integrity of the system as a whole: as the system is designed by and for nerds, and selects for nerds, it is unlikely to change. This function of preservation is necessary in a system that is open to change (to a degree) because its users can change the system as well. While discussing the Visual Editor, I have touched upon the underlying cultural connection between technical prowess and general knowledge and/or intelligence. A further marker that characterises nerd culture is rationalism, and the performance of a specific kind of masculinity, which sees intellectual sparring as an important component of knowledge production (see chapter 1). Talk pages sort contributors on the basis of the ability to

participate in adversarial debate: if one isn't able to bear the aggressive atmosphere of talk pages, then they are not a good fit for the community. As well as this being trivially true – talk pages are an integral part of content production, and not being able to be on talk pages will hamper someone's ability to contribute meaningfully – this is also true under a more general definition of “fit”: if someone doesn't get along with the kind of sparring that goes on in talk pages, then one is not a good fit for Wikipedia as a whole.

Making tools that preserve the culture of the toolmakers is consistent with the framing of Wikipedia's infrastructure as protocol, and can be further illuminated by understanding Wikipedians as a craft community. In doing so, I will also lay the foundations to analyse Wikipedians' coding work in relation to epistemology, in the remainder of this chapter.

Coding communities are craft communities, and in some ways, they operate analogously to other craft communities. Raymond's basic intuition to juxtapose Linux developers and workers building a cathedral was, in essence, sound (Sennett, 2008). Ironically, however, he didn't see the similarities between members of contemporary tech communities and actual stonemasons. Steinmetz (2015) argues that hackers act like guilds: like medieval stonemasons, they are mostly free agents, working for themselves, who come together for specific projects; moreover, they accrue knowledge through apprenticeship, by receiving it from more expert members of the community. Similarly, Wikipedians learn from one another: the vast system of unspoken rules beyond coded policies is learnt informally.

It is important, here, to recall that hackers, free software programmers, and Wikipedians, as well as sharing a broad ethos (see: chapter 1), also share the fact they are, mostly, amateurs – or have started as amateurs, when young, and have become professionals as a result. Not having formal training means both that they depend on others for learning, and that they don't have a certificate that attests their expertise. Not being able to certify expertise is a problem, in terms of quality control: when someone is given a chance to participate to the development of a collective project which, in the case of software as in the case of cathedrals, is critical public infrastructure, there needs to be some measure of trust in their abilities. But, given the relative independence of this kind of craftspeople, that trust cannot always be developed through personal ties. In other words, both medieval stonemasons and amateur programmers need a way of telling whether strangers are experts. In the case of stonemasons, the existence of informal signs of recognition, such as secret handshakes, is documented (Stevenson, 1988). In the case of hackers, cultural practices are used instead, such as trolling (Graham, 2019) and spelling in l33t (Engert, 2009).

Wikipedians find themselves in a similar predicament: since the decentralised editing system doesn't allow for trust to form through closeness, other mechanisms – cultural, such as knowing how to game the system, and technological, such as the sorting technologies described above – are deployed. As argued in section 4.3, sorting technologies don't just tell apart those who are technologically savvy from those who are not. They also test cultural belonging, in the form of the ability to successfully use technologies according to unspoken rules, and, by proxy, they test something like general intelligence, through the assumption that those who are good at programming (nerds) are also likely to succeed in intellectual pursuits more generally.

A case could also be made that something like secret knowledge plays a role here. Stonemasons' knowledge was carefully kept away from the general public, because it allowed to build a lucrative career for oneself, and typically required a form of initiation in order to be obtained (Stevenson,

1988). Similarly, as I explained in section 4.3.2, programming knowledge is regarded as something that is only held by a class of privileged individuals, and programmers do little to dispel this impression. Knowledge of the rules of the game, on Wikipedia, is similarly esoteric (Ford & Geiger, 2012). Esoteric knowledge is also related to programming through the concept of the backdoor – a built-in bypass to a device’s security system, designed to be used by technicians to access the device in cases of emergency; or in other words, to allow to do legally and easily what a hacker does illegally and with ingenuity. It seems significant that l33t, the once-secret, now fairly well-known, spelling common among hackers mentioned above, comes from a program called Back Orifice, developed by a hacker group called Cult of the Dead Cow, which used to access systems through the protocol 31337 – or, in l33t, “e-l-e-e-t”, which happily sounds like “elite” (Engert, 2009). The name “l33t” comes from the 1337 section of the protocol’s name. The shortcuts and subterfuges built into Wikipedia’s policy and design, described above, can be thought of as epistemic backdoors: built-in tunnels through the editing system, accessible only to those who know where to find the entrance.

In sum, the Wikipedia community creates and maintains a network of artefacts that protect its epistemology through protocological means, by leveraging the multistability of technology and cultural markers associated with coding practice. In Galloway’s terms, protocol exerts power in a distributed manner, through desire: Wikipedian sorting technologies, created by the community to further its own aims, also constrain its development and channel power in complex ways, that don’t correspond directly to a bottom-up, technologically mediated imposition of rules. Rather, sorting technologies protect and discipline the community through multistability: the range of uses of any given artefact, corresponding to different users, operates distinctions that serve the wider purpose of affirming and preserving Wikipedia’s epistemic culture. Coding practice allows the creation of protocological, multistable technologies in the first place, as well as providing the markers of belonging that allow those technologies to function.

4.4. The construction of the knowing subject on Wikipedia

As well as to protect their practice, Wikipedians employ protocological means to constitute themselves as knowing subjects, in accordance with community principles, through the operationalisation of specific kinds of expertise. In this section, I will argue that various models of expertise, each characterised by a specific treatment of the identity of the knowing subject, are embedded in Wikipedia’s coded infrastructure. A foundational ideal of the knower, which I have labelled “the clean channel ideal” (Falco, 2023), is embedded in the general setup of Wikipedia’s knowledge-production process; reflecting the principles upheld by Wikipedia’s policy, the clean channel ideal brackets the identity of the knower. A second, antagonistic, model of expertise has emerged over time, aiming to emphasise the identity of the knowing subject, by celebrating positionality. Finally, a third model of expertise, which privileges specific identities, is embedded in those features described above as “sorting technologies”. I will examine how different conceptions of the knower are operationalised, and how they relate to Wikipedia’s policy and culture.

4.4.1. Bracketing the body: the clean channel ideal

As I have argued elsewhere (Falco, 2023) the ideal Wikipedian knowing subject is framed as a

clean channel: abstracted from their own identity, attachments, and feelings. The knower postulated by the clean channel doesn't have a body, in the Merleau-Pontian sense: it is abstracted from location in time and space, history, culture. Or, as we'd say in social science terms, it lacks positionality. The bracketing of the body of the knower is obtained on Wikipedia through various practices, rules and devices, which I will examine in this section.

Symbolic disembodiment is, to an extent, part and parcel of online activity: in the early years of Web 2.0, when Wikipedia was founded, it was commonplace to assume an online identity, separate from one's actual body. The default way of interacting with others was through avatars, or characters that were fictional in name, appearance, often positionality.

In the context of Wikipedia, the first thing to note is that the process of editing, as described in section 4.2.2, is designed to favour impersonal contact. Collaboration is facilitated by the existence of a substrate on which traces can be left for others to pick up: editors modify the Wikipedia page, prompting other editors to check their changes and make changes of their own if necessary. The term "stigmergy" (Elliott, 2006) was later applied to this format of cooperation. Wikipedia can afford an impersonal system of knowledge-production because trust between members of the community is built through the protocological means described in section 4.3 – not by creating relationships based on personal features, but by proving, through action, one's worthiness of being a member of the community. Geek culture, similarly, places emphasis on action: "hackers should be judged by their hacking, not bogus criteria such as degrees, age, race, or position" (Levy, 1994, p. 35).

Additionally, Wikipedia's system of validation, as argued in section 4.2.2, is not rooted in expertise nor experience, at least not in the traditional sense (I will expand on this point in section 4.4.3). Any information that could be acquired by foregrounding the positionality of editors, then, is redundant to the growth and maintenance of the encyclopaedia. Rather, validation is obtained by reducing the influence of body, in a manner that is reminiscent of Ayn Rand's account of knowledge acquisition: as explained in 1.5.3, Rand believed that facts exist in nature, and can be directly apprehended by rationality, when it is unclouded by affect.

Ayn Rand's epistemology, and indeed the clean channel ideal itself, reflect a time-honoured way of framing the scientific subject (Daston & Galison, 2010), as reflected, for instance, in Merton's norms, one of which is disinterestedness (Merton, 1942/1973). Detachment is a guarantee of truth: "scientists aim to distance themselves from the values, vested interests, and emotions generated by their class, race, sex, or unique situation. By decontextualising themselves, they allegedly become detached observers and manipulators of nature." (Hill Collins, 2010, p.255).

Bracketing the body also means hiding personal goals and feelings from view. Disinterestedness and detachment are necessary virtues in an environment where radical collaboration is practiced. It is essential, for a good Wikipedian, not to feel a sense of ownership with respect to one's work. This is important because, first, Wikipedians contribute to the encyclopaedia knowing that anyone else could delete or modify their contribution. Being precious about one's contribution is damaging for the project at large, as it would discourage participation from a large number of people. Secondly, Wikipedia's content is released under "Creative Commons Attribution-ShareAlike 3.0 Unported License and, *unless otherwise noted*, the GNU Free Documentation License" (Wikipedia: Copyrights; original emphasis). As such, it is not owned by its creators. Conversely, no emotional reward is given for editing, nor is it expected. The Wikipedians I interviewed mentioned learning

and the ability to contribute as drivers of their editing. Pleasure was not mentioned.

Downplaying physicality can also be seen as yet another expression of the antisensuality that permeates geek culture. As I started to outline in chapter 3, geek culture has long been characterised by a tendency to ignore, or even harm, the body. There is a tradition, in programming culture, of sacrificing the body to serve the needs of computers, by, for instance, adapting one's circadian rhythms to those of computation, or suffering bad posture as a result of spending long hours at a desk (Turkle, 1984). Computer programmers have also been characterised in mystical terms, as priests or wizards (Ensmenger, 2015), figures similarly associated with a distance from the body, which may manifest as neglect or mortification. On the same continuum, cyberspace has fostered fantasies of immateriality – a world of data, seemingly untethered from material existence, where bodies can be left behind, in a quasi-spiritual fashion (Wertheim, 2000): Wikipedia draws from the same well.

In sum, the system of information management hardwired into Wikipedia at its inception assumes, and to an extent, produces, a bodiless, unperturbed knowing subject, whose main task is to ferry information from a reliable source to Wikipedia's articles. The impersonality of the process aims at minimising interference, emphasising action, and adopting a neutral point of view. This model of knowledge-production is consistent with some of Wikipedia's ideological influences: hacker culture, Ayn Rand, and more broadly, scientific rationalism. In the next section I will describe a direct reaction to the channel model, enacted through artefact that emphasise, rather than conceal, the positionality of editors.

4.4.2. Celebrating bodies: userboxes and the Teahouse

Subjectivity and epistemology are so tightly wound together through Wikipedia's design that coded artefacts built on a conflicting underlying notion of subjectivity are perceived as threats to Wikipedia's epistemology itself. The clean channel ideal has been resisted in multiple ways, through technical means: by coding features that allow space for the expression of personal identity, community building, and the processing of emotions. In the following I will examine two illustrative examples of this practice: userboxes and the Teahouse.

Userboxes are small banners that Wikipedians add to their own user pages to display personal features. They were initially introduced to convey information deemed useful by some: the languages spoken by the user, and the level of editing skill (Westerman, 2009). However, partly because userbox code can be copied, pasted and modified with relatively low technical skill, userboxes evolved over time, beyond displaying encyclopaedia-related, or even factual, information; they became ways in which users expressed themselves. Nowadays, userboxes can be used to display one's adherence to a certain editing philosophy, political affiliation, sexuality, left- or right-handed status, and so on, at various degrees of frivolity.

In 2005, however, userboxes became the centre of a controversy, known as the 'Userbox War' (Westerman, 2009). Those opposing an extensive use of userboxes were worried that displaying one's leanings was a sign of bias. Others that userboxes were actually a tool to work towards neutrality, intended as the sum of a diversity of points of view (*ibid.*). A compromise was reached at the time, but the debate around political userboxes has not died out. If anything, some Wikipedians today seem still unhappy with the very existence userboxes, because of their role as markers of identity. A thread dated September 2020, on the Village Pump (Wikipedia: Village Pump (idea

lab)/Archive 33), one of the community forums, gathers the interventions of many Wikipedians who argue vehemently against userboxes and other kinds of political expression, where the construal of what constitutes ‘political’ expression encroaches more and more on personal matters, such as supporting the LGBT community, or abortion. Some even argue for eliminating “all of such social media paraphernalia from this site and concentrate on building an encyclopedia rather than ‘expressing ourselves’” (Phil Bridger). Others see a benefit in having some kind of way of showing their humanity: “Our userboxes typically are the only outlets for self-expression allowed in Wikipedia. Article-spaces should not reflect our views at all, and talk pages mostly reflect policy-related debates and content disputes.” (Dimadick). The complaints mentioned in this paragraph appeal to Wikipedia’s core mission and values – to build an encyclopaedia based on a neutral point of view – and identify a threat in the use of userboxes: by bringing in the body, userboxes challenge the clean channel model, and indirectly open up avenues for a different, more situated, epistemology.

If userboxes challenge the clean channel model through individual expression, The Teahouse does it by fostering relationships and tending to the emotional lives of editors. The Teahouse, currently, is a forum where Wikipedians discuss their work. It is the spiritual heir – or, according to some, the reincarnation – of an earlier, now closed, forum, the coffee lounge. The coffee lounge was part of Esperanza, a project started in 2005, and closed in 2007, whose aim was to “support the encyclopedia indirectly by encouraging a sense of community” among Wikipedians (Wikipedia: Esperanza). Esperanza offered a support network to an atomised collection of clean channels of information, with practical initiatives supporting mental health (for those affected by “wikistress”), as well as teaching newcomers how Wikipedia works (Wikipedia: Esperanza).

Those who wanted to close the project argued, among other things, that tending to relationships would be a distraction from editing the encyclopaedia (Wikipedia: Esperanza). Some worried that that a small group might become too powerful (see also: the function of piecemeal editing, as described in section 4.2.2), thereby hampering the decentralised, flat power structure of Wikipedia¹². Relationships are then, in this narrative, akin to bias, as they represent an allegiance to a specific group, rather than the community as a whole. Five years after Esperanza disbanded, the Teahouse was created, to address similar needs to those Esperanza was responding to (Wikipedia talk: Esperanza), but only in part: The Teahouse is a space for troubleshooting, and guiding new editors into writing. The archives also function as an FAQ of sorts. Caring mutual support is not common practice.

4.4.3. Sorting through bodies: operationalised trust and positionality

So far, I have examined two models of the Wikipedian knower, that differ in how they frame the bodies of editors, in the extended, Merleau-Pontian sense, which includes their history and context: the clean channel ideal downplays bodies as much as possible, while initiatives like userboxes and the Teahouse celebrate and highlight them. There is a third way in which the interplay between presence and absence of bodies is present in Wikipedia’s design, and that is: as markers of trust, when

¹² It is worth noting that Wikipedia’s power relations are complex, sometimes formally sanctioned, sometimes not (Reagle, 2010; Tkacz, 2015). Esperanza was not a threat to a pre-existing egalitarian system, because there wasn’t one – it was only construed as such.

rules are broken, or sensitive tasks are needed.

In section 4.4.1, I argued that, broadly speaking, the process of editing Wikipedia is designed to do away with trust, and for good reason: in principle, if anyone can edit, then the kind of personal trust that comes from community can't be a part of how the process of knowledge-production is handled. However, the existence of sorting technologies shows that there are – at least – two tiers of imagined editors: the large crowd of casual editors, and a smaller group of Wikipedia-savvy editors, who know how to play the rules of the game, and are entrusted with more sensitive work. The inner circle is defined and maintained through affordance based on specific personal features: technological savvy and belonging to the Wikipedia community.

There are other, more subtle ways in which editors are ranked based on positionality, such as IP addresses. IP addresses are the bare minimum of personal information that has to be provided in order to edit Wikipedia: anonymous editors are identified through IP addresses, in such a distinctive way that “IP address” also means “anonymous editor”. As described above, IP addresses are used by some vandalism detecting tools in order to predict whether an edit might be problematic or not, and consequently, whether it should be flagged up for review by higher-tier editors. Now, an IP address is, in a sense, a literal embodiment of positionality: it allows to estimate the geographical position of an editor. By using geography to sort edits, vandalism tools operationalise trust as a function of positionality. Hence, even among the “second-tier” editors, some are defined, through technology, as deserving of more epistemic trust than others.

One might ask: doesn't this mean that the clean channel ideal, as described above, simply doesn't exist? Clearly parts of Wikipedia's infrastructure uphold it, but others don't. What separates the ideal from resistance and boundary maintenance? I attribute a normative role to the clean channel ideal because, among the notions of knower embedded in Wikipedia's infrastructure, it is the only one that is also supported by the letter of Wikipedia's policies. As outlined in chapter 1, Wikipedia's declared ideological tenets hold that Wikipedia is not written from a specific point of view, and that anyone can edit. Those ideological tenets are consistent with the way Wikipedia's editing system was set up at the beginning, and, as far as I can tell from Larry Sanger's first-hand description of that process, and the analysis of early conversations between Wikipedia's founders-designers, those tenets informed foundational design choices (see Appendix B for details). Further, any attempt at bringing in the bodily dimension of editors, and their positionality, has, historically, been controversial among Wikipedians; conversely, as far as I can tell, no one has ever contested “anyone can edit” as a principle – if anything, criticism has been directed towards the fact that the principle doesn't apply in practice. As far as boundary maintenance is concerned, it is, by definition, contrary to the “anyone can edit” motto, and as such, it is performed in the background, through the design of features that are meant to achieve other goals, such as accuracy and safety.

Despite both situated knower interventions and sorting technologies being in opposition to Wikipedia policy, I have characterised situated knower interventions as antagonistic – hence, a threat to core Wikipedia culture – and sorting technologies, on the contrary, as mechanisms that preserve core Wikipedia culture. The reasons for this choice are, broadly speaking, sociological. First, bringing in positionality in the way userboxes or the Teahouse do would give more epistemic weight to people who are, currently, marginalised within the Wikipedia community, such as new editors (Taraborelli & Ciampaglia, 2010). Secondly, situating technologies contradict Wikipedia's central

epistemic policy, NPOV, in a way that sorting technologies do not. Third, sorting technologies place trust – thus, epistemic influence – in the hands of the same community that generated Wikipedia in the first place: geeks; it is then reasonable to believe that said technology would protect its designers. Finally, the idea that editors can break rules, provided they know how to get away with it, is built into policy itself: the Ignore All Rules policy, together with the argumentative basis of Wikipedia’s community governance (see chapter 1) essentially gives latitude to those who know the rules of the game to behave as they see fit. Sorting technologies act in a similar spirit: they afford experienced Wikipedia the opportunity to exert their influence, while barring newcomers from doing so.

Coded artefacts allow to expose, conceal, and leverage the bodies of Wikipedians, in ways that are either in line, explicitly opposed to, or granting select exception with regards to official policy. In creating and deploying features, Wikipedians affirm their allegiance to one or the other model of the knower, and, consequently, gesture towards distinct approaches to knowledge-production. Thus, Wikipedians express their allegiance to specific epistemologies, turning Wikipedia’s coded infrastructure into a site for epistemic debate, as described in the next section.

4.5. Negotiating epistemology through coding

Wikipedia’s explicit epistemology, defined in policies and guidelines, is not up for debate. In an environment where explicit debate of core epistemological tenets is not culturally appropriate, but contributing through coding is encouraged, programming becomes a way of expressing one’s views on how the production of knowledge should be carried out. In this section, I will argue that designing coded objects, involved in the production of knowledge, which embed values is, in and of itself, a form of debate about epistemology – or in other words, a form of reflexivity. If coding both embeds values, and enacts a form of reflection of those values themselves, then, I will argue, coded artefacts don’t just operationalise, but participate in the development of Wikipedia’s epistemology.

Small ways in which Wikipedians use coding as a means of debate are present in the literature. Geiger (2011) describes pushback against a bot designed to insert signatures where they had been missed. The debate around the bot concealed a wider disagreement against the social norm it reified (signing comments), and was eventually settled through equally technological means, by allowing to opt out (*ibid.*). Forking, as described in section 3.3.1, can also be used as radical technical means of expressing dissent.

Here, I focus on debate around epistemology, operated through technologies that frame the knower in specific ways. In section 4.4, I have described how the clean channel ideal has been put in jeopardy through the creation of features such as userboxes. The activity of creating coded artefacts that embed opposing epistemological views can be seen, I think fruitfully, as a form of debate, or negotiation. Wikipedians are, on one hand, resisting a bedrock of code that they can’t change directly, because under the purview of the Wikimedia Foundation. On the other hand, volunteers are debating among themselves, showing allegiance to opposing sides by adopting features like userboxes, or criticising them; or by modifying them to display certain kinds of information as opposed to others.

As argued in section 4.4, the (discursive) debate around these features is also a veiled debate about the epistemology of Wikipedia itself. Editors are supposed to maintain a neutral attitude

towards their work on Wikipedia: as outlined earlier, sharing one's views is, culturally, frowned upon. Features, however, can be discussed, acting as proxies for underlying epistemologies. For instance, the userbox controversy has sparked a conversation around the relationship between neutrality and identity: to some, marking oneself with userboxes is a breach of neutrality, because it roots editing in a specific point of view (as opposed to a neutral point of view); to others, who conceive of neutrality as the sum of many, diverging, points of view, userboxes are a means of bolstering neutrality.

The discussion on The Village Pump mentioned above (Wikipedia: Village Pump (idea lab)/Archive 33) started by discussing whether userboxes could be divisive, hence potentially harmful to a harmonious collaboration environment, with statements such as:

Delete all political userboxes - Politics is a divisive subject that has no place on Wikipedia when it comes to editor's userpages. Enough of this... we don't need the drama.

(Knowledgekid87)

The conversation then evolves towards the impact that userboxes may have on a given user's editing activity:

We should go back to first principles. What are user pages for? They are not for advocacy, they are there to tell other Wikipedians about yourself, in support of collaborative editing. Ultimately there is a difference between saying "I support X" and saying "Y should not be allowed to exist". [...] Saying you support gay rights is acceptable in that it indicates a possible area of editing interest ("I think you might be interested in X article I just created"), but saying "marriage should be between one man and one woman" is divisive and demeaning to gay Wikipedians.

(Guy)

In-between the many comments dedicated to try and establish exactly what kind of userboxes can be deemed to be political or offensive, other editors bring in the issue of neutrality, wondering whether simply stating one's participation in a specific political group is a neutral statement, and finally touching on the fact that political userboxes, by showing the ideological affinities of a given editor, help identifying them as someone who is trying to push a specific point of view, hence in breach of policy:

I have to say that I do find political userboxes useful, in the same way I find blank edit summaries useful: they're a flag. Often, I'll encounter an editor on a talk page for a charged political topic making a, shall we say, highly strained argument for a particular outcome. I'll initially WP:AGF that they perhaps just have an unusual interpretation of policy, but if I check their user page and it's filled with political boxes, that's useful [...] at the least indicates that they're too blinded by their ideology to look at the issue objectively. I can then choose to

disengage with them or otherwise proceed accordingly. Let's allow POV pushers to out themselves.

(Sdkb)

In practice, a strain of the conversation about userboxes is a conversation on the nature of objectivity: those claiming, like user Guy, that advertising one's perspective can be useful to editing, are essentially assuming that there is something to be gained from having different points of view in the same room. On the other hand, Sdkb is talking about "POV pushers" as editors whose influence should be kept under control. This conversation is also, in part, about authorship: believing that NPOV comes down to personal neutrality means thinking of authorship as, largely, an individual matter. On the other hand, thinking of neutrality as a result of the interaction of positionalities implies understanding Wikipedia's authorship as collective. Positions can, of course, be more nuanced: one might think that objectivity means absence of bias (neutrality in the NPOV sense) or that a higher kind of objectivity is a combination of perspectives (a very different approach from NPOV's).

Another area where coded debate about epistemology is occurring is that of automated anti-vandalism tools: while older tools relied on parameters that, ultimately, created an environment that was hostile to newcomers, ORES is attempting to be more welcoming. It does so by leveraging services that support newcomers, such as the teahouse (Halfaker & Taraborelli, 2015). The designers of ORES also claim to be inspired by feminist ideals (*ibid.*). In a sense, then, insofar as ORES is attempting to serve epistemic justice, it does so by reintroducing the body in Wikipedia's epistemology.

A third example of epistemic debate concerns oral citations. I have described in chapter 3 how a small experiment was carried out, allowing oral citations in articles. It was, however, an isolated occurrence, that had no structural consequences on the way Wikipedia is edited: as of March 2023, all sources of information for Wikipedia articles have to be written (Wikipedia: Verifiability).

In sum, Wikipedia's epistemology is operationalised through coded artefacts in various ways. Expertise is operationalised by leveraging certain markers of identity, skills, and experience. Wikipedia's conception of knowledge is embedded in the mechanisms that operate a selection between what is and is not a suitable contribution, and by offering venues for the sharing of information in formats that are suitable for certain kinds of knowledge, and systems of validation, but not others. By creating, modifying, and adopting features that operationalise specific values, Wikipedians assert their positions on the topic of knowledge-production – negotiating with each other, and with the Wikimedia Foundation. Code, then, as well as embedding an epistemology, is also the means through which an epistemology is created.

4.6. Technoepistemic practices

The discussion so far raises a fundamental question: what is, exactly, the relationship between creating a technology and producing knowledge? In this section, I argue that code constitutes a significant part of Wikipedia's epistemology, and coding practice is a reflexive activity that produces, maintains and challenges it. To be clear, I am not simply stating that Wikipedia's coded

infrastructure reflects the epistemology that is upheld in Wikipedia's policy, and by the members of the Wikipedia community. My claim is stronger: I hold that Wikipedia's epistemology is created, in part, by coding, and is, in part, the result of coding practice. Since Wikipedia's coded infrastructure doesn't necessarily uphold the letter of Wikipedia's policy, I argue that we are not looking at a process of translation, where epistemic values are embedded into technology simpliciter: we are looking at a process of co-creation of an epistemology, which occurs in part through discursive practices, and in part by programming.

As outlined above, toolmaking is also inextricably intertwined with shaping epistemic authority. It defines the knowing subject, and defends the edges of the community of knowers. Wikipedia's infrastructure constructs the knowing subject as disembodied, thereby excluding personal experience, as well as credentials, from the very definition of expertise. Hiding identity, in practice, means excluding certain grounds for expertise, such as personal experience and knowledge accrued outside of Wikipedia. At the same time, sorting technologies give primacy to those who already have technological expertise. Hence, the weight of the contribution of someone who has subject-specific expertise, but isn't technologically savvy, is less than someone who has technical expertise.

By privileging certain kinds of sources, structuring information in a specific way, and certifying what is and is not admissible as true, coded objects contribute to defining the local notion of knowledge. The specific materialities of code (here, loosely defined as text that does things) interact with content in a direct way: the fact that both code and content are text, and sit alongside each other, has epistemic consequences. In other words, Wikipedia's coded infrastructure creates a worldview, and an epistemology, at the same time. The act of coding, then, becomes a way of producing knowledge, and of reflecting on knowledge-production, by designing artefacts that contribute to modifying the knowledge-production process.

Most importantly, the practice of coding itself contributes to Wikipedia's epistemology. I have shown above how Wikipedia's validation system has inherited its shape from programming practice. In doing so, epistemic principles that were applied successfully to code – that is: helped making code that worked, in the context in which code must work – such as radical collaboration and the use of brute force in order to correct mistakes, were applied to the production of knowledge about any possible topic that might fit on Wikipedia.

The epistemic legacy of coding practices is significant. Van der Velden (2013) has remarked how using a versioning system commits the encyclopaedia to present each topic in one way only: the currently displayed article about each topic is the most accurate available account of the topic in question. This ties in with the implicit epistemology of FLOSS development: improving software is a moral imperative, and most importantly, there is such thing as better software – a direction of progress (Leach *et al.*, 2009).

Versioning, I will add, also embeds the idea that knowledge is never final – that it evolves, so that the truth is never achieved once and for all. It should be noted that this is how the free software movement frames code – as something that is never finished. Code is meant to be modified and reused. Programs are, of course, deployed at various stages, and applied to real-world problems, much like Wikipedia can be consulted in its permanently unfinished state, but they are always also starting points for future developers to adapt to their own needs. In FLOSS culture, working with unstable software is a sign that one is doing new, cutting-edge work (Leach *et al.*, 2009). Knowledge-

production, as a result, is understood on Wikipedia as accumulation, through continuous, and unidirectional progress.

The way authorship is framed on Wikipedia also relates to versioning. Coleman (2013) argues that versioning, when used in a programming context, embeds a tension between individualism and collaboration that is typical of FLOSS hackers: on one hand, versioning systems allow to track authorship, by marking each change with the name of the coder who committed it; on the other, the long list of names is a testament to coding as a collective endeavour. In her view, this tallies with a tension found in hacker culture, where hackers routinely work as a collective, but also celebrate individual prowess, often in rather aggressive ways (Coleman, 2013). The corresponding virtues of humility and arrogance are equally celebrated. Similarly, in my own analysis, I have found that Wikipedian culture, while celebrating the communal effort of making the encyclopaedia, is sustained by individuality: editors are expected to take initiative, “be bold”, follow their instinct (see Appendix A for details). Competition is a fundamental mechanism of knowledge-production itself. The Wikipedian process of content creation embeds and upholds this tension. Authors are, in the history page, presented as a whole. The names of editors who contributed, though, are listed – including reverts. Individual edits, and the competition between them, is framed as the real engine of the article.

Even when principles derived from other sources – e.g., openness – inspire the design of coded objects, programmers don’t operationalise them directly into technical features. At least, not entirely: for instance, the principle of “anyone can edit”, once applied to the creation of Wikipedia’s infrastructure, within the concrete context of Wikipedia’s development, was not upheld in the final product. Rather, Wikipedia can be edited by anyone who understands how to use these technologies in a way that is appropriate to the Wikipedia community. When the interface developed with the aim of operationalising Neutral Point of View ends up hiding the names of authors, and evening out their individual voices into an impersonal, bodiless one, the result is the expression of a specific point of view – that of an imagined universal subject, quite clearly situated within a Western rationalist tradition. In both cases, the initial, socially developed principle inspires a technology that, to some extent, betrays it.

Wikipedia’s coded infrastructure embeds its own epistemology, and discussions on said epistemology occur through coding. Again, this is different from saying that epistemic principles found in the wild are translated into artefacts that are then deployed against each other: principles, sometimes, inspire features; other times, features made for a specific, technical, reason (e.g. the Teahouse to encourage newcomers and preserve the very existence of Wikipedia) unwittingly operationalise a specific view of knowledge-production (embedding a model of expertise that is in contrast with a detached, unfeeling, individualistic knower).

There are precedents in framing instruments as embedding knowledge. Baird (2014), argues that instruments embed knowledge about the way they work, as representations (models) and as objects that offer a field of possibilities (as happens, for instance, with measurements). In this sense, technology functions as a bridge between those who use it and the knowledge they produce, and carry with them certain epistemic tenets. According to this account, however, the very act of designing and creating instruments is not an integral part of the knowledge-production process. Baird’s account becomes more useful once combined with Russo (2023)’s notion of techno-scientific

practices, and specifically her use of poiesis. Russo argues that the distinction between the two practices of producing instruments and producing knowledge is, to an extent, illusory: knowledge is a thing that is made, and humans and instruments co-produce knowledge together.

In the context of Wikipedia, the acts of producing instruments to produce knowledge, and producing knowledge with instruments occupy, in my view, the same space. In a technoepistemic practice, producing artefacts, producing knowledge, and reflecting on the production of knowledge are one and the same act. Epistemic aims are achieved through technical means, and technologies are designed based on a certain epistemology. Technical constraints and priorities influence the production of knowledge, while values embedded in, and developed through, toolmaking are irreducible to a discursively expressed epistemology. Knowledge produced through a technoepistemic practice is aggregated and validated through technical means, that constitute in and of themselves grounds for trust in the truthfulness of the content. At the same time, systems of validation are not neutral – rather, they reflect epistemic principles that are shared by those who designed them. Rather than thinking of epistemic and technical practice as informing each other, I suggest we think of them as existing in the same space, and consisting, in part, of the same practices.

Making tools is itself epistemic work, and it can, in some cases, as described above, constitute reflection on epistemology as well. The material process by which a tool is made has an impact on the epistemic values it embeds, and practices of tool production can carry epistemic principles that, then, influence the production of knowledge that is facilitated by the use of said instrument.

4.7. Poietic objectivity

Wikipedia's policy is clear on one point: the criterion for inclusion on Wikipedia is verifiability, not truth (Wikipedia: Verifiability). This, however, doesn't mean that there isn't such thing as truth on Wikipedia; there is something that occupies the same space as truth in knowledge-production process, and that is: the outcome of the process itself – the content of the encyclopaedia. Hence, the definition of truth, on Wikipedia, is the result of the creation and use of coded objects, according to the rules of the community.

Here, and contra previous accounts of Wikipedia's epistemology, I argue that we can think of the knowledge-production process on Wikipedia as a form of poiesis, grounded in the practice of coding. On one hand, Wikipedians make tools, used according to rules that are only known by the members of the community of makers. At the same time, coding practice shapes the subjectivity of Wikipedians, defining the ideal knowing subject and influencing the behaviour of knowers (editors) in practice. Wikipedians and coded objects, in postphenomenological terms, co-constitute each other in relation. Framing Wikipedia's epistemology in this way is consistent with its ideological ties to software development culture, and with a poietic conception of technology, as a means for revealing the truth.

Thus, I argue, knowledge is legitimised by the relation between makers and the artefacts they produce. I referred above to the concept of mechanical objectivity, to describe how Wikipedia validates knowledge, rhetorically, through the – at least, apparent – elimination of the knowing subjects' peculiarities, thereby turning knowers into parts of an automated knowledge-production machine. Here, I suggest a further step is needed to understand Wikipedia, through the notion of

poietic objectivity, a kind of objectivity that leverages toolmaking, and tool use according to the rules of the community of makers, as guarantees of truth.

Throughout this chapter, I have framed Wikipedia editing as a craft. There are obvious commonalities between the work of Wikipedia editors and more traditional forms of craft. Software development, which is the model on which Wikipedia's process is built, has been described as a craft and, as detailed in section 4.3, developers and the hacker community behave much like craft communities. Sennett (2008), building on Plato, argues that the way coding communities work is poietic, in the sense that their making is deeply intertwined with community practice: Linux programmers don't just make software – they make software designed to be reused by others, and commit to sharing knowledge among themselves.

I argue that Wikipedians behave much in the same way: knowing how to code, or how Wikipedia's features work at face value is not sufficient to succeed. Rather, one needs to know the appropriate ways of using and making technologies, to take advantage of Wikipedia's unspoken rules – which are transmitted socially, much like Linux tricks or stonemason techniques – and community structure. Expertise is not defined exclusively in terms of coding, nor just as knowledge of the rules: it is a combination of the two. Expertise, then, is the practice of making within a specific social context – a poiesis. Here, again, I build on Russo (2023), whose account frames scientific knowledge as something that is made. By the same token, if coding the infrastructure regulating the production of knowledge is, as I argued in section 4.6, a way of producing epistemic values, then knowledge isn't the only thing that is made: so is epistemology.

As well as Wikipedians shaping Wikipedia's coded infrastructure, coded artefacts and practice also influence the subjectivity of Wikipedians themselves, as detailed above. As argued in sections 3 and 4, the design Wikipedia's editing process privileges certain ways of being: people with certain characteristics – argumentative, members of the community, technologically savvy, in certain geographical areas, and so forth – are more likely to be Wikipedians in the first place. Further, by encouraging or discouraging certain behaviours through affordance, Wikipedia's infrastructure has, at the very least, the potential to shape the behaviours of editors.

Poietic objectivity, as I define it, is the practice of producing and validating knowledge by making artefacts aimed at supporting the production and validation of knowledge, within the context of, and responding to the ethos of, a specific epistemic community. Makers, within this framework, are granted epistemic authority by virtue of their practice, and, at the same time, submitting themselves to the tools they create. The fact that Wikipedia is built by programmers, and modelled on the development of code, means that coding in itself defines what knowledge is, and how to produce it, but also that knowers are defined by the act of coding. Artefacts, and the act of making itself, shape the subjectivities of makers. Poietic objectivity tallies with the conception of poiesis as revelation: technology is framed as a gateway to the truth, or, in Heideggerian terms, as unconcealment (Heidegger, 1977).

Framing Wikipedia as a programmed encyclopaedia, in the way I have done here, means building on, and moving beyond, accounts centred exclusively around the human element of knowledge-production (see, e.g., Hartelius, 2010 and Fricker, 2012), as well as those, such as Ford's and Van der Velden's, that keep technology into consideration, but don't attribute a central role to it. Ford's account is based on affordances, hence on the way in which Wikipedia's infrastructure is used

by editors, and on the idea that “facts” are accompanied by coded artefacts (Ford, 2022). While of course this is an important part of the picture, it doesn’t capture the depth of the relationship between Wikipedians and coded artefacts, nor does it describe the way in which coding as a practice co-constitutes Wikipedia’s epistemology – and, consequently, the “facts” themselves. Van der Velden’s contribution (2002, 2013), more radically concerned with design, stops at describing how Wikipedia’s infrastructure embeds Western values; here, I grant more agency to the materiality of said infrastructure, and argue that it creates values, in the context of a community of makers.

Coding constitutes Wikipedia’s epistemology materially, ideologically, and through practice. Materially, a system of coded artefacts supports and binds a process of aggregation, validation, and shaping of information. Ideologically, values that became important in software development have inspired Wikipedia’s design, thus affirming themselves as epistemic tenets of the encyclopaedia. The practice of coding, as well as influencing the way in which Wikipedia’s knowledge-production process was designed, also plays an important role in shaping the subjectivities of Wikipedia editors. Wikipedia can be said to function based on an ideal of poietic objectivity, where makers – coders specifically – are given epistemic primacy, in practical terms, as those who create the code that regulates epistemic activity, and symbolically, as the assumed most authoritative members of the community.

Conclusion

In this chapter, I have shown how treating Wikipedians as a coding, as well as an epistemic, community opens up avenues to better understand Wikipedia’s infrastructure. I have argued that Wikipedia’s epistemology is designed in, and negotiated through, coded artefacts. It is impossible, in my view, to understand Wikipedia without recognising how deeply coded artefacts, programming practice, and the culture surrounding it are embedded in its workings. Wikipedia is a programmed encyclopaedia: beyond the trivial acknowledgment that the process of knowledge-production is regulated by coded infrastructure, it is important to recognise the epistemic value of the infrastructure itself. Rather than simply enforcing discursively created policies, coded artefacts shape and enact epistemic principles.

The production of technology, in this case, has to be understood as poiesis – as production in the context of a specific community, stemming from, but not merely mirroring, social epistemic commitments. As a result, culturally savvy programmers are invested with a great deal of authority, which extends beyond the practice of programming itself, to the production of encyclopaedic content. The framework I have started to outline here challenges traditional accounts of the social shaping of technology and theories underlying design for values approaches. In chapter 5, I will outline an account of values in technology based on my analysis of Wikipedia.

5. Towards a material, relational account of the value-ladenness of technology

But while justly respecting great geniuses for their enlightenment, society ought not to degrade the hands by which it is served. [...] how many alleged scholars are there for whom science is in truth only a mechanical art? What real difference is there between a head stuffed with facts without order, without utility, and without connection, and the instinct of an artisan reduced to mechanical operation?

Jean Le Rond d'Alembert, Preliminary Discourse

5.1. Introduction

In chapters 3 and 4, I have shown how specific ideas have influenced the design of, been embedded in, expressed and debated through, coded objects making up en.wikipedia.org. In this chapter, I will collate the values involved in the processes described so far, and draw some theoretical insight about the value-ladenness of technology. I will place empirical observations in dialogue with the established scholarly tradition concerned with the embedding of values in technology. I will show how my analysis of Wikipedia highlights certain gaps in the current understanding of values in technology, and suggest a way forward to expand our understanding of value embeddings.

Here, I argue for complicating existing accounts of the value-ladenness of technology, in three directions. The first direction concerns the subjectivity of the designer: the process of designing impacts the way designers think of themselves, act, see the world. Secondly, we should frame the operationalisation of values into artefacts as a co-construction of the values themselves, making explicit that values before implementation are not the same as values-in-place. Third, we should acknowledge that value embeddings depend, in part, on the internal relations between parts of an object, and on the lineage of similar designs that the object is inserted in. I will conclude by giving some more examples of how these insights can be generalised to other technologies, with special attention to digital technologies.

I will build on postphenomenology and the philosophy of Gilbert Simondon. The two approaches blend well: both are relational accounts, meaning they emphasise the space between players, or actants, and both leverage, intentionally or unintentionally, the notion of affordances. The focus of both philosophies is on the reciprocal entanglements and influences that occur between humans, technology, and the natural world. Simondon's account describes how materiality acts through its physical features, while postphenomenology builds on interaction between objects and humans.

5.2. Wikipedia's value system

In the dissertation so far, I have already dedicated ample space to the description of the specific ways in which Wikipedia's design is value-laden. This section will simply serve as a systematic overview of those values and the connection with their operationalisations, laying the groundwork for the rest of my argument. Based on my research, I have identified five main values informing Wikipedia's design: neutrality, equity, inclusivity, structural openness, and radical collaboration. Given the deep entanglements between hacker ethics, technology, and design, these values can be described as ethical, technical, and aesthetic at the same time, depending on implementation.

5.2.1. Neutrality

Neutrality is, arguably, Wikipedia's foundational value, culturally and technically. It constitutes the basis of policy (section 1.3.2), editor subjectivity (chapter 4, section 4.1) and epistemology (section 4.5). Every interviewee I have asked about values has mentioned it first. Wikipedia's aesthetic and broader approach to design is informed by neutrality: its tone of voice (section 3.3.1), its visual appearance (section 3.6), its ambitions to abstraction from materiality (section 3.4).

5.2.2. Equity

I have listed the Wikimedia Engineering Architecture Principles in chapter 1. Equity is the first one, defined as granting the ability "(to) allow users to consume, create, and interact in a form suitable for their devices, with the connectivity they have, in a language they speak" (Wikimedia Engineering Architecture Principles, 2023). Equity is then broken down in three sub-principles: device, language, and access. The "device" principle refers to the fact that Wikipedia needs to be usable by people with as wide a range of technological arrangements as possible, adapting to a variety of combinations of hardware and connectivity. "Language" is a proxy for geography, and related to the frequently evoked concern around "ubiquity", the ambition for Wikipedia content to be available everywhere. Finally "access" here refers to accessibility, which traditionally, in Web design, is understood as the value that aims at levelling the playing field when it comes to disabilities (Kennedy, 2012). For a website to be accessible, it means that its content is available to people with a variety of sensory abilities, including those who have, for instance, a visual impairment. Early conversations between Wikipedia developers about the appearance of the default skin (see section 4.4) can also be read as conversations about accessibility, as they bring in the ability of the reader to customise their reading experience based on their specific needs.

5.2.3. Inclusivity

Inclusivity is a fundamental principle of Wikipedia, as testified by the motto "anyone can edit". It could be argued that the most radical innovation brought about by Wikipedia is not that the content of the encyclopaedia is available to a large group of people, but that producing content is. A lot has been said about what "anyone" means. The original tagline was written by a group of people with two main concerns in mind: access levels, and credentials. In line with hacker culture, Wikipedia doesn't lock doors nor require formal education – it was designed without technical barriers to participation and allowing space for self-taught editors to contribute. Measures limiting access

started to be included in 2005, after the “Siegenthaler incident”, in which a hoax article about journalist John Siegenthaler was published (Reagle, 2010).

Over time, however, the notion of inclusivity has become more nuanced – starting with the notion of accessibility, above – and come to include other factors that may preclude or discourage certain groups of people from participating to the project. One such factor is technical skill: as shown in chapter 4, Wikipedia’s interface can be disheartening to those who are unfamiliar with markup languages. Objects such as the Visual Editor have been created to extend inclusivity in this sense. The other prominent factor is demographics: Wikipedia’s editorship is rather uniform. As described in chapter 4, whether the identity of the editors matters is a contested point. The debate, as argued in section 4.4, is partly carried out through design interventions such as userboxes and oral citations.

5.2.4. Structural openness

Structural openness is, primarily, a technical value. Wikipedia is designed with the aim of making interventions possible. As I have described in chapter 4, Wikipedians contribute by editing content and programming; openness in this sense applies to the material fabric of Wikipedia, which allows for relatively direct, minimally mediated intervention. Wikipedia is structurally open from a social point of view as well: the community has multiple direct channels of communication with leadership, who is expected to respond; the social aspect of openness is implemented in facilities such as forums and user talk pages belonging to employees of the Wikimedia Foundation, up to the CEO.

5.2.5. Radical collaboration

As explained in chapter 1, Wikipedia’s process is based on the ideal of radical collaboration: anyone can edit anyone’s work. Radical collaboration, in and of itself, is a practice: it consists of a particular way of working collectively on the same artefact. I include it here as a value because it is ethically charged, and because it correlates with certain virtues that are valued within the community. Radical collaboration is ethically charged because it entails a form of selflessness, of nobility: giving up one’s contribution for editing means relinquishing, at least to an extent, recognition. Of course, Wikipedians do get rewarded for their work, on occasion; such occurrences are, however, rare. Most Wikipedians edit in anonymity, and, as argued in chapter 4, their subjectivity is partly shaped by the absence of a reward. The emphasis on humility often found in Wikipedian discourse – Wikipedians have been labelled “janitors of knowledge” (Sundin, 2011), administrators are said to “take the mop” when they are elected – is a result of radical collaboration. Radical collaboration is also implemented throughout the editing process, by allowing editing other people’s work in the first place, which is arguably the foundation of the whole knowledge-production enterprise.

Wikipedia’s values are operationalised in various ways through technology. Sometimes the embedding occurs implicitly, as is the case for neutrality and radical collaboration, which were native components of the worldview of Wikipedia’s founders – from Ayn Rand and hacker culture, respectively. In other cases, values are articulated explicitly as desiderata for engineers to keep into account when developing the platform, as is the case for equity and inclusivity. Either way, what interests me is the relationship between values and artefacts: how is operationalisation achieved? What does it mean for values to be embedded in artefacts? In the next section, I will explain why I

find existing accounts of value-ladenness in technology unsatisfactory on this point, by applying some concepts from the philosophy of Gilbert Simondon, and specifically leveraging his critique of hylomorphism. Section 3, together with section 2, will then set the stage for a proposal (section 5.5, 5.6 and 5.7) based on my own observations of Wikipedia (section 5.4).

5.3. Values in technology and hylomorphism

In this section, I will argue that existing theories of value-ladenness, within the fields of philosophy of technology, STS, and sociology of technology, ignore, or downplay to a fault, the role of materiality in the design process. In this section, I will give a broad outline of these approaches, and show how, when it comes to the nature and definition of the values embedded in technology, they lend primacy to human actors, ignoring or downplaying the agency of artefacts. I will argue, borrowing the language used by philosopher Gilbert Simondon, that these approaches adopt a hylomorphic approach to technology: that is, they assume that human action shapes artefacts, independently of the nature of the artefacts themselves. Having laid out some conceptual challenges to hylomorphism, I will proceed to examine empirical evidence against it.

The value-ladenness of technology is explored by a large body of work. The literature investigating the relationship between value-led design can be roughly divided in two broad areas: critical approaches, rooted in science and technology studies or sociology of technology, and work rooted in philosophy of technology. Within philosophy of technology, two distinct, but complementary – and, indeed, sometimes fruitfully brought together – traditions exist: postphenomenology, and a cluster of theories that I will label, collectively, as “value-led design”.

The first group focuses on the web of social relations around artefacts, and the social shaping of technology (including, among others: Eubanks, 2018; Latour, 1990, 2000, 2005; Lemonnier, 1993; Pinch & Bijker, 1984; Wajcman & MacKenzie, 1999; Winner, 1980). Theories in this space do two main things: they describe how society influences the way technology is designed, and hence embeds values that are prevalent in society, and how the use or deployment of certain artefacts have social consequences. Langdon Winner describes artefacts as the reification of social relations, through material means; at the same time, technology has an impact on society: it transforms social relations (Winner, 1980). A canonic example from Winner’s work is the viaducts built under the supervision of Robert Moses, in New York, with the aim of reducing the presence of Black people in certain areas: the viaducts were too low to let busses through, and the vast majority of Black newyorkers didn’t own a car (*ibid.*). Theories of this nature tend not to linger on the mechanisms by which embedding certain views in artefacts occurs, focusing mostly on the demographics of those who design technologies, and on the social consequences of building and deploying such artefacts.

Within this group, Bruno Latour, is the author who most closely looks at how embedding works. His Actor-Network Theory (ANT) defines the nexus between society and technology in terms of translations and statements (Latour, 1990, 2005). In *Technology is society made durable* (1990), Latour describes the case of a hotelier attaching a heavy keyring to room keys, in order to ensure that said keys are returned to the desk when the guest leaves the building, instead of being carried around. According to Latour, the hefty keyring incorporates an imperative (‘leave your keys at the front desk’); the same imperative could be written out on a sign, but it would sit differently within the Web

of social relations surrounding the key. A sign inviting customers to return the keyring is, for instance, easier to ignore than a heavy weight in one's pocket; hence, if a guest is particularly inclined to keep the key, it will be easier for them to do so, thereby shifting power towards the customer and away from the hotelier.

I have already introduced postphenomenology and value-led design theories, in chapter 2. The role of materiality in postphenomenology is relatively important: the artefact constitutes one end of an affordance-based process of co-constitution. I will expand more thoroughly in section 5.5, but for now it's important to note that, from the point of view of my present concern – the role of materiality within the design process – postphenomenology is silent, for the good reason that it tends to apply to the context of use, not the context of design. Value-led design theories, on the other hand, implicitly follow the same pattern as the sociological accounts described above, in that values, defined independently of the envisaged artefact, then influence the design of the final product. I will speak more to this in the remainder of the chapter.

Materiality has a role in these theories, but a limited one. Viaducts and key ring act thanks to their material properties. Viaducts physically impede the passage of buses. Latour's account hinges on materiality: incorporating 'leave the keys' into the weight and size of the keyring translates the order, rather than transmitting it, as a sign would. Materiality also changes the imperative hotel guests are obeying: "Customers no longer leave their room keys; instead, they get rid of an unwieldy object that deforms their pockets" (p. 105). Similarly, postphenomenology postulates that objects have a form of agency, that expresses itself through affordance.

The crucial problem with these accounts is not the absence of materiality, but its function of enforcer of human agendas. The overarching mechanism by which these accounts work is an assumption that – sometimes ideally, sometimes in practice – social values and relations are designed into an artefact, which receives them as embeddings, or is shaped by them. Winner speaks explicitly of "intentions embodied in physical form" (Winner, 1980, p. 125). The ethical, and/or the social, are the primary drivers of the process: they are deployed on matter that is assumed to be, broadly speaking, inert.

It is useful, here, to leverage the concept of hylomorphism, as applied to technology by Simondon. As I have explained in section 2.4.2, Simondon notes how the relationship between a mould and the material being moulded, often conceptualised as a one-way process where the mould is active, and the material is passive – that is, the mould shapes the material – is, rather, one of mutual shaping: the mould, itself an object that has been made, is built to impose a shape on the specific kind of matter it is made for. In Simondon's case, such matter is clay – one could easily see how moulds made for other materials, such as, for instance, jelly or molten metal, would need to be made differently.

Simondon's analogy can be applied, in my view, to the understanding of the design process implicit in the theories sketched above – casting humans as shapers, and artefacts as being shaped. When we apply the concept of hylomorphism to some of the theories sketched above, it becomes clear how they are, indeed, hylomorphic, and how hylomorphism may constitute an issue in terms of their explanatory and prescriptive power, by erasing the role of materiality, and specifically the importance of, borrowing again from Simondon, implicit forms. The significance of this absence will become clear over the course of this chapter. Sociological theories of technology – social shaping of

technology, ANT – explicitly assign a significantly primary role to social forces. Similarly, value-led design theories, while foregrounding values defined within the context of design, leaving societal influence further in the background, still attribute absolute moral primacy to human decision-making. In both cases, ignoring or, in my view, excessively downplaying the role of material constraints on the development of new technologies.

Analysing technology through a hylomorphic lens forcibly obscures the role of materiality in the design process: the approaches described above simply do not offer the tools to understand it. The existence of these conceptual issues is bolstered by empirical work. My research on Wikipedia provides evidence for an even more articulated network of influences; in the next section, I will list some case studies, all examined over the course of the dissertation so far, that can be leveraged to extend the role of materiality in value-led design.

5.4. Challenges to the hylomorphic model, from my work on Wikipedia

As well as being conceptually unconvincing, the hylomorphic model can't explain some of the phenomena I encountered in my exploration of Wikipedia's design. As I will show below, social processes only partially explain why Wikipedia is designed the way it is, and the values embedded in its design are related to, but don't overlap exactly with, the ideals expressed in policy. Wikipedia's coded infrastructure itself plays a significant role in defining its own ethics, by participating in the construction of the subjectivities of its designers, and of the values embedded in its design; by passively enacting trade-offs through conflicting implementations; by directing its own evolution; and by harbouring inherited values in its functioning. This section will constitute the basis to generalise further, in sections 5, 6, and 7, where I will pave the way for an understanding of values in technologies that concerns itself more deeply with materiality.

5.4.1. Materiality and the subjectivity of designers

In chapter 4, I have highlighted some ways in which a certain kind of self – detached, disinterested, “civil” – is constructed through the editing process. Here I will describe three mechanisms which are, in my view, involved in the process of construction: selection, technologies of the self, and practice.

Selection is operated through affordance: as I have argued in section 4.3, Wikipedia features objects that I have termed “sorting technologies”, as they separate editors in categories, based on whether, and to what extent, they are practically able to edit Wikipedia. Some of these technologies assume users to have certain skills, are built to cater to a specific kind of knower, and, as a result, contribute to attracting users with a specific profile. The effect of this initial sorting is magnified by the editing interface, which assigns highest levels of epistemic credibility – and, consequently, of contributing privileges – to tech-savvy users (section 3.2). Other technologies leverage positionality more directly: vandalism tools, for instance, may filter out editors based on their geographic location (section 3.1). On the whole, the basic threshold that needs to be cleared in order to edit Wikipedia is rather high, and can be discouraging for newcomers (section 3.4).

The second mechanism that impacts the subjectivity of editors pertains to self-policing, relying

on a mechanism that can be subsumed, in my view, to Foucault's technologies of the self (see 2.4 for an overview). As discussed, in order to participate meaningfully to the project, Wikipedians have to adopt certain attitudes, such as civility, as defined within Wikipedian culture (section 1.3.2), detachment from their own contributions (sections 4.4), and specific ways of expressing anger (section 1.3.2). Being part of Wikipedia's editing process means disciplining one's feelings and behaviours in specific, locally defined ways. Participating to a materially regimented process, in other words, has an impact on subjectivity. In the case of Wikipedia, those who use tools to edit are also, to some extent, the designers of the tools themselves. Hence the subjectivity of designers is co-shaped by the tools they design themselves. Technologies of the self are mediated and bounded by technologies of production created by the same people whose subjectivity is being moulded.

The editing process offers a good example of how Wikipedians shape themselves through editing. Operating under a regime of radical collaboration – the practice of allowing anyone to edit anyone's contribution – entails an ideal of selflessness and self-abnegation (see section 4.4). At the same time, though, since such an ideal is unattainable by most, radical collaboration fosters an atmosphere of continuous conflict, where elaborate means of battling are deployed by editors who are not inclined to defer to others. Radical collaboration, then, both contributes to upholding a certain ideal of detachment, and to its failure in practice.

As a result, means of dispute resolution have to be put in place. However, due to cultural and policy reasons (see chapter 1), dispute resolution is a complex game, likely to be won by those who are familiar with its subtleties. It has been noted elsewhere (Ford & Geiger, 2012) that Wikipedia editors become successful by learning how to play by the rules, and fitting within Wikipedia's complex cultural landscape. This requires work on the self: speaking the language of neutrality in the terms defined by Wikipedian culture, becoming more and more detached from one's edits, managing one's feelings when editing (see chapter 1 for behaviour policies). A certain level of performed detachment, approximating, but not fulfilling, the ideal entailed by radical collaboration, is then necessary in order to thrive.

Finally, editors learn through making. In order to edit Wikipedia, one has to learn the process, first, but also ends up knowing something about the content of the pages one edits, and, if one decides to contribute to coding objects such as bots or gadgets, or help out with the core code, one acquires coding skills too. An indication that Wikipedia coders improve over time is that documentation pages exist that introduce Wikipedians to basic concepts, with the implicit expectation that more can be learnt. Encouraging people to learn and improve is consistent with FOSS ethos: free sharing of knowledge, and helping people develop skills is part of the movement's historical mission. At the time Wikipedia was founded, hacklabs – environments where people could come together and learn how to code, manage IT infrastructure, and commune with other people with similar interests – were popping up in Europe and the US. Deep personal involvement through learning is a tenet of Wikipedians' activity.

5.4.2. Material affordances and value embeddings

In this section, I aim to show how the material characteristics of Wikipedia have a bearing on value embeddings, by lending itself to absorb certain kinds of values, and in a certain form. Wikipedia's purpose – collective knowledge-production – and material qualities define a specific

subset of implementable values, and contributes to defining those values themselves, based on what is possible through the available infrastructure.

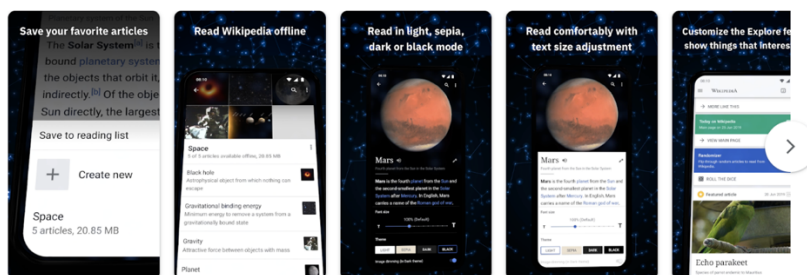
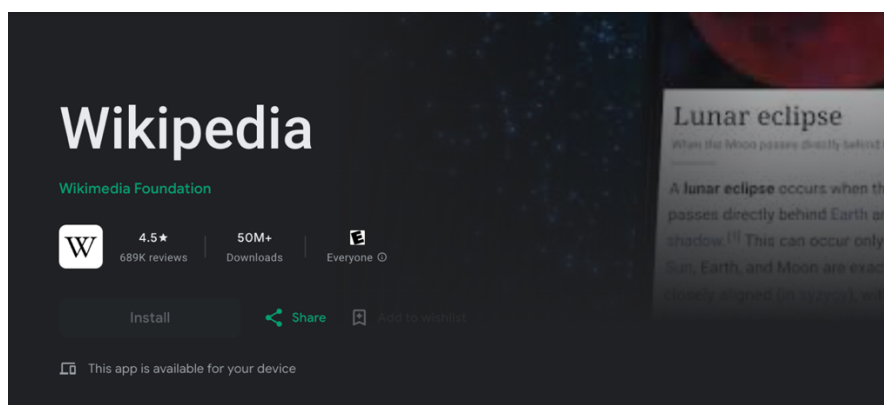
Wikipedia's website, designed to produce an encyclopaedia, lends itself, primarily, to embedding epistemic values. Additionally, due to the collective nature of Wikipedia as a project, political and social values can be designed into Wikipedia's infrastructure as well. Crucially, however, Wikipedia cannot contain any and all values. Filial piety, for instance, is not something that Wikipedia, as an artefact, can support or undermine, because it does not mediate relationships between relatives, nor contain information around family ties. Hence, filial piety is simply beyond its jurisdiction. Wikipedia's implementation, its purpose, and the idiosyncratic structures created to fulfil its purpose, have a bearing on the kind of values that can be embedded in it. A rough qualitative correspondence, then, exists between values and artefact.

Secondly, Wikipedia's coded infrastructure partly defines the values themselves, through operationalisation. I have already examined how proposals for reform of Wikipedia put forward by Heather Ford and Maja Van der Velden operationalise different conceptions of justice, in relation to both their respective ethical commitments and definition of epistemic justice, and the different features of the website they consider to be salient (section 3.3.3). Here, I will focus on how equity is defined through Wikipedia's infrastructure as it currently is. As described in section 5.2.2, equity is defined on Wikipedia as a value that concerns geography and accessibility (in the disability sense). A first observation can be made here, that this definition of equity is already influenced by context: equity is defined differently in, for instance, education (Kaur, 2012), health (Peter, 2001), administration (Hart, 1974), economics (Thurow, 1981). Partly¹³ because these are engineering principles pertaining to a website, as opposed to, say, radiology equipment, they focus on what can reasonably, easily, and communally, be addressed through digital features. For instance, accessibility can be obtained in various ways, through a minimal, and customisable, interface (see chapter 3). Equity is, in part, defined as accessibility, because accessibility is implementable on a website.

Most importantly, the way equity is defined in principle changes as the process gets closer to operationalisation. Equity, conceptualised as applying to reading and editing Wikipedia, ends up being limited to reading only: strategy papers building on this principle, and the technical measures taken to enforce it, are mostly geared around consuming Wikipedia's content. Wikimedia Foundation documents discuss using a variety of platforms, making content available in various geographical areas, and on a variety of devices. The emphasis is, throughout, on audience. In terms of format and platforms, Wikipedia content is imagined as being available regardless of physical substrate (see section 3.3.1). Geography is mostly operationalised as coverage: the priority, across Wikipedia initiatives, is to expand the coverage of currently underrepresented areas, sometimes in ways that don't necessarily amplify local voices, or support local epistemologies (section 3.3.3). Cross-platform content is, as of now, mostly imagined; geographical expansion, mostly obtained through social means, such as local projects. Closer to my interests here, and crucially, already widespread, is the use of mobile devices for editing.

¹³ Ideological influences, as mentioned in the rest of the thesis, also play a role here. I do not intend to swerve towards determinism in these last few pages. Here, however, my concern is to foreground the complementary role of materiality.

Mobile usage is an equity issue: mobile connectivity is growing worldwide, particularly in developing countries (Ceci, 2024). Hence, if the concern with equity sits at the intersection of increasing editorship in underrepresented communities, connectivity, and devices, it is important for Wikipedia to improve its presence on mobile. Reading Wikipedia on mobile is relatively simple: it can be done through a browser, or various apps, including the official Wikipedia app, published by the Wikimedia Foundation. Editing on mobile, on the other hand, is more difficult, for various reasons. First, there is an anatomical barrier: mobiles are smaller than computers, don't have a keyboard or a mouse, making typing long strings of text, including markup, challenging. Mobile screens don't lend themselves to carry out research that requires cross-referencing, which entails having multiple windows or tabs open at once, copying and pasting content. Precision in highlighting, selecting, and changing text is also hard to achieve, and made more complicated by the intervention of autofill. Because of these limitations, mobile devices are typically used for ancillary work, such as making small edits on the fly, before sitting down to do the real work (Wikipedia: Editing on mobile devices). An already challenging hardware setup is made even more unsuitable for editing by the fact that the official Wikipedia app doesn't allow to use Visual Editor – only the default one, where precision, for instance in editing tags, is even more crucial – and doesn't present an option to access talk pages rapidly (Wikipedia: Mobile communication bugs). The Wikipedia app, in terms of affordances, is then an app for reading, rather than editing, Wikipedia, incurring in issues of epistemic asymmetry I have described throughout chapter 3 and 4. Additionally, and perhaps as a result of the obvious technical limitations of editing on mobile, the Wikipedia app is marketed as an app for reading: its description doesn't mention editing facilities (Fig. 5.1), nor do the introductory screens presented after downloading the app (Fig. 5.2).



About this app →

The best Wikipedia experience on your Mobile device. Ad-free and free of charge, forever. With the official Wikipedia app, you can search and explore 40+ million articles in 300+ languages, no matter where you are.

Figure 5.1 Wikipedia app as presented on Google Play, the Android app store.




Figure 5.2 Introductory screens to the official Wikipedia app.

In Fig. 5.2, the emphasis is clearly on exploration; considerations around equity clearly underlie the description provided: Wikipedia can be read “no matter where you are”, in “300+ languages”, and options for customising the reading interface are presented. Editing is never mentioned. In the context of the material arrangements of connectivity, devices, geography, interface design, and human anatomy, equity acquires a narrower meaning: its definition ceases, to a meaningful extent, to apply to producing knowledge – clashing with Wikipedia’s most fundamental cultural value, access to knowledge-production – shifting heavily towards consumption.

5.4.3. Material trade-offs

Value trade-offs are a significant object of study in value-led design scholarship. However, attention is given primarily to trade-offs that happen outside of the object itself. Trade-offs are discussed as something that needs to be attended to before the object is created – where conflicts between values due to be embedded into future objects are to be solved before the objects are designed and produced – or after it has left the factory floor and is out in the world being used, at which point users, who were not involved in the design of the object, might bring in their own values and interests, or unforeseen ethical consequences might show themselves as a result of the interaction between the object and the outside world. What, then, should we make of those trade-offs that emerge within the object, and are one step removed from the action of the humans who make and use it?

Wikipedia presents such a case. In section 4.3.3, I have described the case in which trying to make it easier to edit markup, through the Visual Editor, ended up creating a situation where editors whose native language is written right-to-left are penalised. This trade-off is a function of the fact that code is text, and that the way Wikipedia is edited mixes text that says things (content), and text that does things (code). In other words, the issue does not reside in the value of inclusivity, which inspired the creation of the various interacting artefacts in the first place, but the relations between them, and the double functionality of text in the editing interface.

5.4.4. Value change

A similar question emerges in the realm of value change. There is an acknowledgment, in the literature, that objects evolve over time, in various ways: obsolescence, deterioration, update, change in use, re-purposing, maintenance, and so forth. The literature on value change and technology acknowledges that values designed into an artefact can change over time due to, for instance, changes in the moral landscape surrounding the object in question, impacting on the relative importance of values, shifts in their definition or specification (Van De Poel, 2021). So there is some benefit in building in the ability to cope with changes in the environment: Van de Poel (2021) mentions adaptability (the ability of an artefact to change, which increases with modularity), flexibility (a range of different uses for the same design) and robustness (ability to respect values in a range of different instances of the same product). Change over time can also be leveraged in and of itself to embed a certain value. For instance, Verbeek talks about objects made by a company called Eternally Yours, that are designed to age beautifully: the design changes over time, in a pleasant way, thereby coupling ageing and beauty in the eye of the users, and making it less likely for the objects to be thrown away (Verbeek, 2005).

However, as I have described in section 1.5.1, Wikipedia was set up to be coded collectively, in a decentralised manner, and as a result, is now difficult to modify and, in some ways, to understand. That ethos of transparency and openness inspired Wikipedia in the first place has, over time, and without any change of heart from the community, turned into closeness and elitism. Of course, some human activity is involved here: volunteers have added lines of code to Wikipedia's code base. But the real action, the change in ethos, occurred through a shift in relations between parts of the monolith, hence internal to the monolith: relations became entangled and calcified in such a way that they are hard to manipulate. That internal relations are the main issue here is supported by the way the Principal Software Engineer, Moriel Schottlender, talks about it as well. As she explained to me in conversation, she is trying to move Wikipedia to a decoupled system, that is, a system that is modular, where different parts do different things – hence, the relations between parts are looser – with the explicit ultimate aim of making the system more manageable. Secondly, the affect around discussions of the monolith is telling: there is a sense of frustration, and impotence – a tacit acknowledgment that human action has limited purchase in this situation. The object is asserting a degree of autonomy, in the sense that it cannot easily be changed by humans. This level of autonomy is not acknowledged, nor investigated, in most literature about value-ladenness in technology.

5.4.5. Material legacy

Another aspect of Wikipedia that needs to be considered is how much it has inherited from previous technologies and processes. I have made the case earlier (section 4.2.2) that the process by which Wikipedia is edited is a form of versioning – a process inherited from traditional programming, which allows for distributed, parallel work. Versioning carries certain values, such as self-determination and autonomy: by fostering direct access to code, thereby eliminating filters and priority, versioning allows coders to work together in an unmediated and parallel, hence non-hierarchical, way. Similarly, wiki technology itself embeds a non-elitist ethos, because it affords relatively easy intervention by a large group of people, through a commonly available, and relatively

familiar, tool (a Web browser). Hence, there is a continuity between the values embedded in versioning and wiki technology, and those close to Wikipedians' hearts. One might argue that perhaps those technologies were chosen precisely for that reason: it is most certainly true of wiki technology, which in and of itself affords decentralised, non-hierarchical collaboration, as acknowledged by Larry Sanger himself in his memoir (Sanger, 2005). It is not necessarily true of versioning, however, which might have been adopted simply out of cultural habit – it was common in the environment Wikipedia founders belonged to – but brought with it, unawares, affordances that embedded specific values.

There are cases in which the tools used embed values that contradict stated Wikipedia goals. In section 4.2.3, I have described STiki, a now-obsolete vandalism detection tool that is nonetheless interesting because of its ethical implications, and because the team who programmed it detailed their thinking behind the design in a paper (West *et al.*, 2010). One of the interesting aspects of STiki's functioning is that bypasses certain Wikipedia core principles: STiki works by ignoring the content of potentially disruptive edits and only looking at metadata – that is, data about the edit, containing information about time, place, and so forth. Metadata is then evaluated based on reputation, which, among other things, depends on the seniority of the editor who made the edit, and their country of origin. STiki also allows for human intervention: it alerts administrators, who are then tasked to take action. Now, the ethos of Wikipedia is, explicitly, that anyone can edit, anonymous edits are allowed and encouraged, the place where someone is editing from shouldn't matter, and the selection of edits is non-hierarchical. So why use reputation, geography, and human intervention to filter out edits? The creators of STiki defend this choice as highly effective, but don't compare it to other methods. What they do say, though, is that their tool was based on spam filters – which, at the time when STiki was made (2010), used reputation and geography to evaluate whether an email could be spam.

I have already discussed the ethical implications of using these specific parameters in chapter 4. Here I want to focus the attention on how those ethical pitfalls are not the result of Wikipedia's ethos being turned into a feature of the platform – but of the decision to use, for practical reasons, an existing mechanism, developed in a different context, to solve a Wikipedia problem. Using a method based on spam filters caused a disruption in the embedding of values into Wikipedia's infrastructure, by essentially infecting it with the values underlying the original tool. Essentially, this means that, depending on the case, re-using an existing technology can either reinforce or go against the dominant ethos of the Wikipedia community. One can also say that seniority and geography, in practice, do matter (again: see sorting technologies), even if they are not the mainstream, or even ideal, way in which expertise is determined. And indeed I have described how competing ideals of expertise can be coded at the same time, in different parts of the platform (section 4.5). If this is the case, then, we can say that STiki takes a side in Wikipedia's coded debate, and that the side it takes is, at least in part, the result of its technical history.

In this section I have shown how materiality plays a role in shaping the value-ladenness of Wikipedia's coded infrastructure, in two ways. For one thing, materiality has a bearing on aspects of the design process that have been traditionally considered fundamental to understanding value-ladenness in technology – the identity of designers and the definition of values. At the same time, Wikipedia's infrastructure applies pressure to the design process through internal features – by

allowing or resisting the embedding of specific values, imposing trade-offs, driving value change over time, and carrying the moral legacy of its predecessors. In the next sections, I will draw paths to generalise the Wikipedia-specific insights above, with the ultimate ambition of expanding our understanding of value-ladenness.

5.5. Co-constitution of subject and object in the design process

Thus far, I have aimed to make two broad points: that existing theories of value-ladenness in technology lack tools to address materiality, and that my work on Wikipedia shows that those tools are needed. If we are to make sense of artefacts like Wikipedia – complex, digital, sociotechnical systems, extensively investigated at the time of writing – we need to start building such tools. In the following, I outline some directions for development, by building on existing work in philosophy of technology and STS and defining three concepts that are, in my view, central to this project: co-constitution of subject and object in the design process, co-construction of values, and material dynamics.

In this section and, to an extent, in the next, I will use the concept of co-constitution, borrowed from postphenomenology, to understand the relation between designers and the objects they create. Co-constitution relies on the notions of affordance and mediation: objects lend themselves to certain uses, while at the same time interjecting and reshaping the user's lifeworld. Thinking of artefacts and designers as co-constituting each other places them in a relational space where subjectivity and artefact respond seamlessly to each other, through the practice of making.

Verbeek defines co-constitution as the phenomenon by which an object (the artefact) and subject (the user) define each other in relation:

Mediation does not simply take place between a subject and an object, but rather co-shapes subjectivity and objectivity.

(Verbeek, 2005, p. 130)

As I argued in section 4.4, Wikipedians' subjectivity is shaped by the act of editing Wikipedia, at the same time as Wikipedia, the object, is shaped by the activity of editors and volunteers. Putting it this way may create the impression that I am talking about a loop: Wikipedians shape Wikipedia which shapes Wikipedians. However, taking a relational approach means moving beyond a simple causal loop: subject and object exist in relation – to an extent, they exist in function of each other. A person is not a Wikipediaian without Wikipedia: their identity, as a person, is partly defined by their activity of making and unmaking Wikipedia; their identity as a Wikipediaian is entirely defined by their activity on Wikipedia. Wikipedia couldn't exist without Wikipedians, and vice versa.

There is plenty of evidence that design practices influence the subjectivities of designers (see below for some examples). As far as I can tell, this co-shaping occurs at two levels: a personal level, concerning habits, skills and knowledge, and a collective level, which involves meaning making, political action, and cultural traits that end up defining communities of practice grown around design and making. The development of a distinctive geek subjectivity was made possible by, and is inextricably tied to, the ability to explore and build afforded by Web 2.0 (Bell, 2009) as well as the

lifestyle required by programming (Ensmenger, 2015; Turkle, 1984). The formative role of programming is beautifully explained in the following passage, written by a member of the Italian hacker collective Autistici & Inventati:

This was my idea, when I started doing A/I. A place like a dojo, a gym, where you find instruments, to pool experiences together. And perhaps even more than that, a place where, by virtue of the fact that there are other people around you using the same tools, sharing occurs almost through osmosis, even if you don't think about it directly.

Starting from a structure which, in essence, does logistic and strategic work in the field of communications, you end up with something that fosters a cultural shift and collective growth. Not by imposing meaning from above, but by letting meaning emerge from the whole thing.

(Ginox, in Beritelli, 2012, p. 245, my translation)

The formative role of the design process is not exclusive to programming and occurs in a range of ways. In participatory contexts, the design process can lead to personal development (Østergaard *et al.*, 2018; Lundmark, 2018, involving nurses and counsellors, respectively), fostering a sense of community (Taylor *et al.*, 2016); creating political awareness (Costanza-Chock, 2020). It is important to note that the interaction with objects is crucial in these processes of transformation, and goes hand in hand with technologies of the self:

This role [of genuine participants] is established – for example – when users are not just answering questions in an interview about their point of view or knowledge of a particular issue, but are asked to step up, take the pen in hand, stand in front of the large whiteboard together with fellow colleagues and designers, and participate in drawing and sketching how the work process unfolds as seen from their perspectives. Inviting users to such collective discussions and reflections requires a trustful and confiding relationship between all participants. Any user needs to participate willingly as a way of working both as themselves (respecting their individual and group's/community's genuine interests) and with themselves (being concentrated present in order to sense how they feel about an issue, being open towards reflections on their own opinions) as well as for the task and the project (contributing to the achievement of the shared and agreed-upon goals of the design task and design project at hand)

(Jesper Simonsen & Robertson, 2013, p.5)

It's also worth remembering that participatory design processes are aimed at developing tools that participants themselves are going to use, by leveraging knowledge acquired by the participants themselves through practice. The subjectivity of participants is entangled in multiple ways with the design process: subjectivities and materiality, designers and artefacts, co-constitute each other in relation through the design process.

When I point towards a relation between designers making tools and tools making designers, I

am not simply talking about the fact that we make things that then shape us (a point that has been abundantly made, and within postphenomenology itself – see Ihde & Malafouris, 2019, for a fundamental contribution). The loop involving us making something and the something making us is a rather wide one: the object is created, exists in the world, and enters relations with humans, who are shaped themselves. The relationship I describe here is much more intimate, and created simultaneously between those involved: it occurs during the design process itself, before the artefact is even finalised or produced. Subject and object are placed in the same relational space, as described by co-constitution in postphenomenology (Verbeek, 2005), or by the concept of enactive individuation, which highlights the continued, iterative engagement between artefact, body and affect of the designers during the design process (Poulsen, 2019).

While it has been established that subjectivities are built and shaped in relation to design processes, the role of co-constitution in the context of design has not been formalised yet. In my view, when we think about embedding values in tech, we need to think about how the people doing the embedding are changed by the process, and what that, in turn, does to the artefact. The relational quality of design needs to be foregrounded more explicitly. In the context of value-led design, understanding the relation between designer and object as one of co-constitution allows to be mindful of the co-evolution of designer and object, and to conceptualise values as something that emerges through the design process, rather than being established in advance, and imposed onto the artefact.

5.6. Co-construction of values

Earlier, in section 5.4.2, I argued that Wikipedia doesn't embed values by absorbing them from human actors as they are – rather, values are embedded in a specific version, or definition, which is the result of interaction between human actors and the platform. Here, I will argue that a similar process occurs at the point of design: values, in order to be embodied through artefacts, have to be operationalised, turned into design requirements that fit the artefact being designed. This is acknowledged in the literature: Latour speaks of translation, value-led design of a process of specification or operationalisation.

The critical question I am asking here is: what is the role of artefacts in the process of operationalisation? As I have argued in section 5.3, current accounts assume an asymmetrical, hylomorphic process: in sociological terms, values that are important to a community – or even just to the part of the community who is in charge of designing or commissioning technological advances – influence what kind of values become embedded in technology. Value-led design, while more granular in describing how values are embedded in objects, also gives primacy to human actors: the selection of values is postulated to occur among humans. Human actors, be it designers or potential users, decide on what values they want to embed, and embed them.

In the following, I will argue that artefacts influence value embeddings via two processes: one of selection, whereby only certain values can be reasonably embedded in an object, and one of construction, meaning that abstract values, as they are conceived of by humans, are re-defined in a way that is compatible with operationalisation in a specific artefact.

5.6.1. Selection

In section 5.4, I noted how not all values can be reasonably embedded on Wikipedia. The purpose and material existence of Wikipedia is such that it lends itself to certain embeddings, and not others. To expand on this, I will look at some examples of simpler technologies where the divide between what can and cannot be embedded comes more readily to the fore.

Two classic cases of value-laden technologies, analysed in the literature, are the speedbump and anti-homeless benches. In Latour's analysis, the speedbump, by its very presence, enforces an ethical rule ("slow down") through material threat ("if you don't slow down, your car might get damaged") (Latour, 2000). Anti-homeless benches feature some raised slats or other vertical elements that interrupt the length of the seat, to impede lying down, and, consequently, sleep (Rosenberger, 2020). The former embeds an imperative along the lines of "slow down", with its implied value of road safety; the latter, the imperative, valid only for a certain subgroup of people, "don't sleep here", which implies a form of elitism and disdain of poverty. Speedbumps and anti-homeless benches live in specific social and material landscapes: speedbumps in a context where car suspensions break if you drive fast over an obstacle (and are expensive to fix), and anti-homeless benches in a context where sleeping outside under duress is considered socially undesirable. Within their respective contexts, the structure and placement of the objects themselves makes it possible for them to do what they do. A speedbump wouldn't work if it were made of foam. And neither of these objects could, within existing social arrangements, embed "truth" or "efficiency", as their function doesn't concern areas where those values would be relevant. The material affordances of speedbumps and benches, within the context of their design and use, limit their ethical jurisdiction.

As a more complex example, whose process of value embedding is thoroughly documented, I will use an experimental game called RAPUNSEL (Flanagan *et al.*, 2008), designed with the express aim of teaching programming, maths, and science skills to teenage girls, with particular attention to those from marginalised backgrounds. The embedding of values in this case is intentional, planned, and described in a paper – the description of the explicit practices around value-embeddings makes this game a great object of analysis for my purposes.

Broadly speaking, RAPUNSEL was designed to promote social justice; values were elicited from stakeholders (*ibid.*). Because of the nature of the game, a lot of the values implemented have to do with learning, and leverage the reward system of the videogame. Hence, the main values operationalised have to do, ultimately, with skills: encouraging a variety of skills, promoting the learning of certain specific skills (maths, programming, science). The game also promotes secondary values such as cooperation, by rewarding teamwork. Other values are included as values that are typically associated with digital technologies, such as privacy and transparency. The purpose of the videogame (learning) and its materiality (being a digital object), intersecting with cultural meanings associated with digital technologies (widespread worries around privacy) contribute to selecting the values that matter the most. At the ground level, artefacts influence their own value embeddings through selection: their purpose and physical features define the range of values they can operationalise.

5.6.2. Construction

Values are not generally understood to be features of objects. In order to work as specifications (i.e. properties of objects), values must be turned into radically different entities. Additionally, the design of an object is, in and of itself, technologically mediated. Translation is, itself, mediated by (material) tools (Ihde, 1990). In modern design processes, values are manipulated using various kinds of objects, such as mock-ups, preliminary designs, images.

The construction of values is relational: it is borne out of the encounter between the – wider, human, abstract – value-in-the-wild and the materiality of the object in question. In the case of Wikipedia and universality, for instance, universality is translated into a requirement that can reasonably be implemented in the infrastructure of an online encyclopaedia. Or, in simondonian terms, according to the implicit forms inherent in websites. Thinking of Wikipedia in terms of implicit forms allows to provide an elegant description of the coded dispute about competing epistemologies described in section 4.5. As I explained in section 5.4.2, the core of Wikipedia's code is not accessible, for sociotechnical reasons: hence, the original values Wikipedia was founded on, and that are embedded in the editing process, are impossible to change. Bespoke code, however, offers an opportunity for intervention; so does the creation of forums. In the dispute between objectivity framed as neutrality, and objectivity framed as highlighting positionality, those in the former camp have Wikipedia's main infrastructure on side. Those in the latter intervene how they can, by creating artefacts such as userboxes and forums; in turn, only certain forms of situated knowledge, and specifically those rooted in expressing positionality, can be upheld through programming, while proposals such as Van der Velden's, which would entail a complete overhaul of the platform, are unimplementable (see chapter 4). Similarly, objectivity as neutrality can be implemented through text, by distributing writing among a group of people, and therefore achieving the appearance of a neutral point of view (section 3.3.1), but applying the same strategy to video contents creates the opposite effect: it highlights the various points of view of the contributors, thereby getting closer to a perspectival approach to objectivity. While the principle is the same, and Wikimedia Foundation documents seem to assume the effect is as well, the materiality of the medium entails a significant difference in result: video cannot conceal perspective in the same way as text does – rather, it amplifies it.

In RAPUNSEL's case, material affordances co-defined diversity. The authors describe how diversity was conceptualised based on the markers of race and gender, but neither is described as being operationalised through the design of the game. The one kind of diversity whose operationalisation is described in depth is diversity in terms of learning styles and skills, which was translated into a variety of ways in which players could get positive feedback throughout the game: players could get points for home decoration, dancing, and other activities, as well as fulfilling the main goal of the game. The idea of diversity adopted by the engineers, then, was partly defined on the basis of what was possible, and what seemed appropriate: through designing a learning game, diversity was defined in terms of learning.

Something similar happened with social justice. In RAPUNSEL, the way social justice is turned into specifications is through a series of mechanisms that make learning STEM more appealing to teenage girls. Now, social justice as a whole cannot be reduced to teenage girls finding STEM a bit

more fun; it cannot even be reduced to increasing the representation of women in STEM. And indeed, this is not what the designers working on RAPUNSEL, and the rest of the team, are trying to do. But they are constructing a specific version of social justice that fits the material constraints of the object they are making – its purpose and affordances.

It is important to note, as well, that co-construction itself is materially mediated, because the design process is. The idea of introducing diversity as a value occurred to the participants while evaluating technical issues. The way values were constructed with the participation of potential users was also mediated by artefacts such as mock-ups and prototypes. This is how the authors describe part of the process:

Noting the pleasure users derived from building and dressing up characters and from manipulating them in the game to engage in relationships with other characters via flirting, dancing, and other social behaviors, RAPUNSEL designers inferred users' valuation of creative self-expression, authorship, community, and collaboration.

(Flanagan et al., p. 337)

In this case, values are constructed through the interaction of users, designers and objects, carrying their own material affordances and implicit forms. Using mock-ups and prototypes is a common approach in value sensitive design (Friedman & Hendry, 2019). As a side note, the passage above also highlights the role of affective (“pleasure”) and subjectivity-building (“self-expression, authorship, community, and collaboration”) components described in section 5.5.

Further, we can think of construction as a result of underdetermination: once values are defined, trade-offs are resolved, and so forth, what ultimately is kept into consideration when approaching the metaphorical drawing board is not strictly specified, and may be decided based on contingent circumstances (Van de Poel, 2013), including feasibility. In other words, designers create blueprints within the constraints of what is doable. And what is doable is dictated by the material affordances of artefacts.

The role of artefacts in value-ladenness doesn't stop at the design stage. Artefacts are dynamic entities, that change over time, present internal relations, and influence the design of future objects. In the next section, I will outline an account of how the internal organisation of complex objects can result in unexpected moral outcomes, and argue that we should pay attention to the lineage of objects, as well as the history of the context of their creation.

5.7. Material dynamics

The morally salient features of artefacts – both in value-led design theories and postphenomenology – are not static. Artefacts are conceived as dynamic, at least to an extent: in some cases, the passing of time itself can be leveraged to express values. In what follows I will expand this line of reasoning, arguing that there are further ways in which we can think about artefacts as dynamic, and incorporate them into a broader approach to value-led design. First, I will address the issue of trade-offs, specifically how complex objects may contain structural elements that create

trade-offs between values, thereby mediating trade-offs through materiality. Secondly, I will look at change generated by internal pressures, as opposed to the environmental pressures typically associated with value change. Finally, I will take a step back and look at artefacts as part of a lineage, and ask how we can develop a sufficiently granular account of inheritance of values when it comes to technology.

5.7.1. Internal dynamics

Trade-offs between values are robustly investigated within the social side of design – not so, to the best of my knowledge, when it comes to the material realm, leaving an important aspect of value-led design largely unexamined. As I have outlined in section 5.4.3, once something has been designed based on certain values, and is supposed to have embedded them, those values end up interacting, through the materiality of the object, in ways that sometimes change the meaning of the original set.

The internal moral turmoil complex objects present is absent from the philosophy of technology literature, which focuses mainly on the question of trade-offs between human values, mainly as they occur before implementation: if a group of stakeholders hold competing values, a solution needs to be found in order to prioritise values and, consequently, specifications (Miller *et al.*, 2007; Van de Kaa *et al.*, 2020). Solving the question should lead to implementing one end of the trade-off, or finding ways to reconcile them. Trade-offs can also occur at the point of use: a feature that was developed in response to a certain value specification can, as a result of contextual changes, end up promoting the development of different, conflicting values (Van De Poel, 2021). For instance, the use of wearable technologies in the workplace, designed to encourage workers to collaborate, with a view to improve productivity and wellbeing, can also end up commodifying wellbeing (Iliadis & Pedersen, 2018).

Material trade-offs, internal to artefacts, are documented as ethical conundrums – but not conceptualised as embeddings. A well-investigated example is the case of the trade-off between transparency and accuracy in AI (Cooper *et al.*, 2021; Wang, 2023). The values themselves, as one might describe them in ethical terms, are not inherently in opposition: one can reasonably value both transparency and accuracy (while one couldn't, say, reasonably value both racism and anti-racism). There is no in-principle trade-off there. The trade-off occurs through implementation, and because of the specific object that is being designed. Sometimes such trade-offs are not foreseen, so no debate is had, but they are still present in existing artefacts. Materiality, then, mediates internal trade-offs between values, in a way that is partly independent from human action.

Another case in which the general phenomenon is acknowledged, but internal dynamics are ignored, is value change. In the value-led design literature on value change, change is always somehow directed by humans: the role of the object is to allow for certain changes to happen, or respond to external, human, changes. Here, I want to bring attention to internal change, that is not directed by humans, and yet, still consequential.

It is helpful to go back to Wikipedia. As described in section 5.4.3, Wikipedia's website was built to afford change, and a specific kind of change at that: it was built to afford participatory, self-directed, decentralised coding. However, the very structurelessness of the design process meant that Wikipedia's codebase has grown organically, creating a complex, and to an extent unknowable,

system of relations, which in turn, twenty years later, makes change very difficult, thereby hampering those values of open participation that were upheld in the first place. Wikipedia's monolith is hard to change because the internal relations between its parts are so complicated that it is very difficult, sometimes impossible, to know the consequences of a certain kind of change. Other times, the consequences are known, undesirable, and difficult to prevent. I say this is a relational problem for good reason: in a decoupled system – that is, a modular system, where each part is relatively more autonomous from the rest than in a monolith – change is much easier.

The inner lives of objects have a moral quality: they may uphold, ignore, or override human value embeddings. This can be understood by observing specific objects at a specific point in time, and by tracking how the structure of an object changes over time, in part through engagement with the world, in part due to its implicit forms.

5.7.2. Lineage

Existing accounts of technology tend to bracket the material lineage of artefacts, or keep it into account in a way that does not concern itself with details. There is of course a rich tradition, in sociology and history of technology, of analysing the lineage of artefacts, but not at the level of granularity that value-led design dedicates to values in technology: what is missing, from historically informed analyses of technology, is a study of material mechanisms by which certain values are carried forward, and others aren't. The history of AI, for instance, is problematically intertwined with the racist history of intelligence, which in turn has important social repercussions, in terms of imaginaries, diversity in the industry, and inequities (Benjamin, 2019a, 2019b; Cave, 2020). To the best of my knowledge, however, it hasn't been convincingly argued, yet, that machine learning as a technology is inherently racist. For such an argument to be possible – and I believe it is – the granularity of the analysis would need to reach down to mechanisms; to explain exactly what material features of machine learning have an exclusionary quality, and, using a relational approach, in what contexts. Viaducts were racist, in Langdon Winner's account, because inserted in a context of systemic inequality which meant white people were driving cars, and Black people riding buses. They are not, however, universally racist.

In light of the absences above, I suggest investigating the lineage of objects by blending the historical, context-sensitive approach of sociologists, and the granular lens adopted in value-led design: which implicit forms have been inherited, and have carried meaning or values with them? Which have been dropped? If the value-ladenness of an object was due to the interaction between two features of the object, are they both still present in current iterations of that object? With code, the notion of lineage can be extended a step beyond improvement, because, partly thanks to the open source movement, code can, and often is, reused. Coding is bricolage. Consequently, if a piece of software is designed with certain values in mind, and is then, entirely or in part, re-used elsewhere, its presence in the new artefact will be morally relevant.

This comes to the fore more clearly if we think about the fact that code is made of solutions to problems, at different scales (see STiki, described above): a certain problem is solved in a certain way by a certain piece of code; when that piece of code is then re-used, the same solution is applied to a different context. As well as carrying archaeological values with it, that specific solution can interact with other parts of the new coded object, each carrying their own ideological baggage. As well as a

continuity at the level of designs, then, there is a material continuity in coding that is relevant to the operation of embedding values in technology. As argued in section 5.4.5, lineage doesn't necessarily mean that what was true of earlier versions of an artefact is true of its descendants; in the case of code, it doesn't mean that inherited code brings with it embedded values wholesale. It does however mean that we need to pay attention to how inherited features interact with new features or features inherited from other places. Again, it's a case of adopting a granular way of looking at technology, of the kind adopted in value-led design, and combine it with attention to historical and material detail, mostly found in sociological and historical literature.

It should be acknowledged that digital technologies are a privileged standpoint to look at internal dynamics and lineage in the way I have described above. Because of the ease of expanding and the necessity of updating software, the internal relations between parts of coded objects can shift often. Coding, as a practice, encourages borrowing, for both cultural and material reasons: code affords copying and modifying, because it is text, displayed on a rewritable medium. At the same time, however, designs for three-dimensional objects, such as furniture, houses, clothing, tools and so forth, have historically been inherited and modified as well. Hence the principle of a technological lineage, and what it entails, can potentially be applied to a wide range of artefacts.

In this section I have outlined what it would mean, for a theory of value-ladenness, to look at objects as dynamic entities, featuring complex internal interactions, and influencing the design, and consequently, value-ladenness, of other objects. Material dynamics are easily visible in digital objects, due to their structure and the design culture producing them; code is not, however, a special case: just a paradigmatic one. Similar conclusions can be drawn about other kinds of objects.

Conclusion

In this chapter, I have argued that values are not simply embedded into objects – they are shaped in the process of making, with the participation of various forms of materiality: the artefact itself, the tools employed in the design process, the raw matter that constitutes the object itself.

I started by sketching a critique of existing accounts of value-ladenness in technology based on the concept of hylomorphism, as applied to technology by Gilbert Simondon. I have argued that existing theories tend to be hylomorphic: they assume a design process where a form, be it values, or social processes, is imposed on matter. The role of objects is mostly limited to that of enforcers of human values, whether those of the designers, or, when misuse or creative use occurs, the users. Here, I have argued that such accounts, by dismissing the role of materiality, are missing part of the picture.

Finally, I have imagined the contours of a theoretical proposal that takes seriously the role of materiality in value-ladenness. The proposal is geared around three main concepts: the co-constitution of designer and object, the co-construction of values, and material dynamics. First, we should start by framing the design process in relational terms, as a space in which designers and the artefact they design co-constitute each other in relation. Secondly, we should complicate the notion of embedding in a similar vein, by acknowledging the reciprocal, materially mediated relations that exist between human actors who bring values to the table and the affordances of the artefact being created, as well as of the various objects that are involved in the design of the final artefact. Finally, we should attend to the relations between materialities, in terms of internal relations between the

various parts of specific technologies, and by situating technologies within a lineage of artefacts.

Conclusion

The theoretical upshot of my work, in relation to the research I conducted on Wikipedia, is described extensively in chapter 5. I won't repeat myself here – I will simply pull some threads together, to describe the narrative arc of the thesis as a whole.

Over the course of this dissertation, I have shown how the infrastructure that constitutes Wikipedia, made out of various connected digital artefacts, does more than embedding values. It co-creates them – on one hand, by welcoming or resisting intervention, and by being a site of ideological negotiation; on the other hand, by suggesting, implicitly, what values are important, what constitutes a moral good in the first place. Beyond affording intervention to humans – being a substrate or a tool for ethical and epistemic meaning making – Wikipedia's platform offers up technical values to be turned into social, epistemic, aesthetic values. This is true, for instance, of forkability: as I have shown, forkability in its original formulation informed design because of its practical advantages, concerning safety and distribution of code. Forkability then, through the community that created Wikipedia, became an epistemic value as well.

Programming practice is a key component of Wikipedia's culture, a core influence on its design: the backbone of the project's lived – as opposed to defined in theory – epistemology. As I have argued in my analysis, the concrete circumstances in which coders have worked define the way Wikipedia produces knowledge. Hence the remark often flippantly thrown around in Wikipedian environments – “Wikipedia works in practice, but not in theory” – assumes a deeper meaning: explanations of Wikipedia's functioning are not to be found in theory at all. Practice is epistemology.

A side-effect of the partial overlap between programming and creating Wikipedia's content is that the flavour of Wikipedia's community matches cultural traits of hacker culture. The effect of this phenomenon is two-fold. First, Wikipedia inherited assumptions found in hacker culture – downplaying the role of the body, faith in machinery, anti-aesthetic leanings, connecting intelligence and skill with the ability to code. Secondly, because of the connection between taste and belonging to specific communities, Wikipedia can be an uncomfortable space for those who don't participate in hacker culture.

Both in terms of aesthetics – both narrowly construed as matters of taste, and broadly construed as considerations around experience – and technological affordances, Wikipedia, then, betrays its mandate of being an encyclopaedia anyone can edit. The consequences are glaring: the largest, most-used reference tool in the world is written by and about a minority of people. It shouldn't be a surprise that a purportedly universal epistemic project is not indeed universal; it is a staple of Science and Technology Studies that all knowledge as situated. Here I highlight some of the mechanisms situating an epistemology, some of which have not, to the best of my knowledge, been described before.

The broader horizon of my work is to give dignity to artefacts themselves, to what they do when humans aren't looking. Since we insist on creating more and more complex machines, we need to make peace with our inability to fully own them. As we attempt to reckon with their power, it is becoming increasingly clear that they are not just our extensions. Machines are, to an extent, a

mystery: an unknown that requires us to work in order to be understood, and will never wholly reveal itself to us, because it evolves continuously. We can't catch up. Appropriately, those who understand the relational nature of technology move from hermeneutics, a practice borne out of the analysis of sacred texts.

My ambition for this study is to show what can be gained from taking materiality seriously, without forgetting about the social and cultural context of design. Future research might address other artefacts in a similar vein, while creating new tools to further our understanding of value embeddings.

As we interrogate new digital artefacts and their meaning – at the time of writing, artificial intelligence is at the centre of everyone's attention; who knows what will be next – we need more than critique. Saying where artefacts come from doesn't cut it anymore, because we are moving further and further away from their origins, witnessing more and more iterations of the same designs (and questioning what it even means for two objects to be the same); describing the context of their production, and assuming a direct infusion of values from social to material, is not sufficient either, as artefacts push back, charm, question, scare, evolve away from, their designers. We need solid tools to understand change, if we want to make sound decisions on what to do with it. Once again, then, back to the things themselves.

Bibliography

- Abramowski, N. (2022). So, is HTML a Programming Language? Retrieved from: <https://careerfoundry.com/en/blog/web-development/is-html-a-programming-language/>
- Ainsworth, T. (2024). Form vs. Matter, The Stanford Encyclopaedia of Philosophy (Fall 2024 Edition). Retrieved from: <https://plato.stanford.edu/archives/fall2024/entries/form-matter/>
- Ambrosio, C. (2014). Objectivity and Visual Practices in Science and Art. In M. C. Galavotti, D. Dieks, W. J. Gonzalez, S. Hartmann, T. Uebel, & M. Weber (Eds.), *New Directions in the Philosophy of Science* (pp. 353–368). Springer International Publishing. https://doi.org/10.1007/978-3-319-04382-1_24
- Ankerson, M. S. (2018). Dot-Com Design: The Rise of a Usable, Social, Commercial Web. NYU Press.
- Antin, J., Cheshire, C., & Nov, O. (2012). Technology-mediated contributions: Editing behaviors among new wikipedians. [conference paper]. CSCW '12. Seattle, Washington, US.
- Ashford, D. (2014). “A New Concept of Egoism”: The Late Modernism of Ayn Rand. *Modernism/Modernity*, 21(4), 977–995. <https://doi.org/10.1353/mod.2014.0084>
- Ayers, P., Matthews, C., & Yates, B. (2008). *How Wikipedia Works*. No Starch Press.
- Bacon, F. (1620/1909–14). Proœmium of the Instauration Magna, in *Famous Prefaces*. The Harvard Classics.

- Bauer, W. (1965). The encyclopaedia in China. *Journal of World History*. 9(1), 665. Retrieved from <https://www.proquest.com/scholarly-journals/encyclopaedia-china/docview/1298905882/se-2>
- Baytiyeh, H., & Pfaffman, J. (2009). Why be a Wikipedian. *Proceedings of the 9th International Conference on Computer Supported Collaborative Learning - CSCL'09*, 1, 434–443. <https://doi.org/10.3115/1600053.1600117>
- Bear, J. B., & Collier, B. (2016). Where are the Women in Wikipedia? Understanding the Different Psychological Experiences of Men and Women in Wikipedia. *Sex Roles*, 74(5–6), 254–265. <https://doi.org/10.1007/s11199-015-0573-y>
- Bednar, K., & Spiekermann, S. (2022). Eliciting Values for Technology Design with Moral Philosophy: An Empirical Exploration of Effects and Shortcomings. *Science, Technology, & Human Values*, 01622439221122595. <https://doi.org/10.1177/01622439221122595>
- Bell, D. (2009). Geek Myths: Technomasculinities in Cybercultures. In Biricik, A., & Hearn, J. (Eds.), *Proceedings from GEXcel Theme 2: Deconstructing the Hegemony of Men and Masculinities Conference 27–29 April 2009*.
- Benjamin, R. (2019b). *Race after Technology: Abolitionist Tools for the New Jim Code*. Polity.
- Benjamin, R. (Ed.). (2019a). *Captivating Technology: Race, Carceral Technoscience, and Liberatory Imagination in Everyday Life*. Duke University Press.
- Beritelli, L. (Ed.). (2012). *+kaos – 10 anni di hacking e mediattivismo*. xBook. <https://www.autistici.org/static/img/book/ai-book-kaos.pdf>

- Biricik, A., & Hearn, J. (Eds.). (2009). *Proceedings from GEXcel Theme 2: Deconstructing the Hegemony of Men and Masculinities Conference 27–29 April 2009*.
- Blair, A. (2006). A Europeanist's Perspective. In *Organizing Knowledge: Encyclopaedic Activities in the Pre-Eighteenth Century Islamic World*, Gerhard Endress (ed.). Brill.
- Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. MIT Press.
- boyd, d. (2008). How Can Qualitative Internet Researchers Define the Boundaries of Their Projects: A Response to Christine Hine, in A. Markham & N. Baym (Eds.), *Internet Inquiry: Conversations about Method*. Sage.
- Broughton, J. (2008). *Wikipedia: The Missing Manual*. O'Reilly.
- Cave, S. (2020). The Problem with Intelligence: Its Value-Laden History and the Future of AI. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 29–35.
<https://doi.org/10.1145/3375627.3375813>
- Ceci, L. (2024, February 2). *Mobile Internet usage worldwide—Statistics & Facts*.
<https://www.statista.com/topics/779/mobile-Internet/#topicOverview>
- Chabot, P. (2013). *The Philosophy of Simondon Between Technology and Individuation*. Bloomsbury.
- Coleman, E. G. (2013). *Coding freedom: The ethics and aesthetics of hacking*. Princeton University Press.
- Collins, P. H. (2000). *Black Feminist Thought*. Routledge.

Collison, R. L. (1964). *Encyclopaedias: Their history throughout the ages*. Hafner Pub Co.

<https://archive.org/details/encyclopaediasth0000coll>

Cooper, A. F., Abrams, E., & Na, N. (2021). Emergent Unfairness in Algorithmic Fairness-Accuracy

Trade-Off Research. *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and*

Society, 46–54. <https://doi.org/10.1145/3461702.3462519>

Costanza-Chock, S. (2020). *Design Justice: Community-Led Practices to Build the Worlds We Need*.

Knowledge Unlatched.

Cram, J. (2020). UNMASK - Digital Thinking with Brutalism. ISEA Symposium.

Daston, L. and Galison, P. (2010). *Objectivity*. Princeton University Press.

Daub, A. (2020). *What tech calls thinking*. Farrar, Straus and Giroux.

de Boer, B., Te Molder, H., & Verbeek, P.-P. (2018). The Perspective of the Instruments: Mediating

Collectivity. *Foundations of Science*, 23(4), 739–755. [https://doi.org/10.1007/s10699-018-](https://doi.org/10.1007/s10699-018-9545-3)

[9545-3](https://doi.org/10.1007/s10699-018-9545-3)

DeClute, D. (2022). Is HTML a programming language? Retrieved from:

<https://www.theserverside.com/feature/Is-HTML-a-programming-language>

Dogramaci, S. (2015). Communist Conventions for Deductive Reasoning. *Noûs*, 49(4), 776–799.

<https://doi.org/10.1111/nous.12025>

Elliott, M. (2006). Stigmergic Collaboration: The Evolution of Group Work: Introduction. *M/C*

Journal, 9(2). <https://doi.org/10.5204/mcj.2599>

Engert, C. A. (2009). Are you with me? How online conversants make 133t connections.

Thresholds, California State University, Stanislaus.

<https://www.csustan.edu/sites/default/files/honors/documents/journals/thresholds/Engert.pdf>

Ensmenger, N. (2015). "Beards, Sandals, and Other Signs of Rugged Individualism": Masculine culture within the computing professions. *Osiris*, 30.

Eubanks, V. (2018). *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St Martin's Press.

Falco, E. (2023). How to Read Wikipedia: Design Choices and the Knowing Subject, in Lerg, C. A., Östling, J., & Weiß, J. (Eds.). *History of Intellectual Culture 2/2023: Modes of Publication*. De Gruyter. <https://doi.org/10.1515/9783111078038>

Fallis, D. (2008). Toward an epistemology of Wikipedia. *Journal of the American Society for Information Science and Technology*, 59(10), 1662–1674.
<https://doi.org/10.1002/asi.20870>

Feenberg, A. L. (2017). Concretizing Simondon and Constructivism: A Recursive Contribution to the Theory of Concretization. *Science, Technology, & Human Values*, 42(1), 62–85.
<https://doi.org/10.1177/0162243916661763>

Felker, K. (2002). Ariadne's thread: Hypertext, writing, and the World Wide Web. *Library Hi Tech*, 20(3), 325–339. <https://doi-org.libproxy.ucl.ac.uk/10.1108/07378830210444522>

Flanagan, M., Howe, D. C., & Nissembaum, H. (2008). Embodying Values in Technology. In J. Van Den Hoven & J. Weckert (Eds.), *Information Technology and Moral Philosophy*.

Ford, H. (2022). *Writing the revolution*. MIT Press.

Ford, H., & Geiger, R. S. (2012). “Writing up rather than writing down”: Becoming Wikipedia Literate. WikiSym '12, Linz, Austria.

Ford, H., & Wajcman, J. (2017). ‘Anyone can edit’, not everyone does: Wikipedia’s infrastructure and the gender gap. *Social Studies of Science*, 47(4), 511–527.
<https://doi.org/10.1177/0306312717692172>

Foucault, M. (1969). *L’archéologie du savoir*. Paris: Gallimard.

Foucault, M. (1994). *Ethics: Subjectivity and Truth* (P. Rabineau, Ed.). The New Press.

Free Software Foundation (2019). The Free Encyclopedia Project.
<https://www.gnu.org/encyclopedia/encyclopedia.en.html>

Free Software Foundation (2024). What is Free Software? <https://www.gnu.org/philosophy/free-sw.en.html>

Fricker, M. (2012). Group Testimony? The Making of A Collective Good Informant. *Philosophy and Phenomenological Research*, 84(2), 249–276. <https://doi.org/10.1111/j.1933-1592.2011.00565.x>

Friedenthal, A. J. (2017). CITATION NEEDED *Wikipedia and the Mutability of the Past* (pp. 145–158). University Press of Mississippi.

Friedman, B., & Hendry, D. (2019). *Value sensitive design*. The MIT Press.

Frost-Arnold (2018). Wikipedia. In Coady, D., & Chase, J. (eds.). *The Routledge Handbook of Applied Epistemology*. Routledge. <https://doi.org/10.4324/9781315679099>

- Fuster Morell, M. (2011). The Wikimedia Foundation and the governance of Wikipedia's infrastructure, in Lovink, G., & Tkacz, N. (Eds.), *Critical point of view: A Wikipedia reader*. Inst. of Network Cultures.
- Galloway, A. R. (2004). *Protocol: How control exists after decentralization*. MIT Press.
- Geiger, R. S. (2011). In Lovink, G., & Tkacz, N. (Eds.). (2011). *Critical point of view: A Wikipedia reader*. Inst. of Network Cultures.
- Geiger, R. S. (2014). Bots, bespoke, code and the materiality of software platforms. *Information, Communication & Society*, 17(3), 342–356.
<https://doi.org/10.1080/1369118X.2013.873069>
- Graham, E. (2019). Boundary maintenance and the origins of trolling. *New Media & Society*, 21(9), 2029–2047. <https://doi.org/10.1177/1461444819837561>
- Grune, D. (1986). *Concurrent Versions System, A Method for Independent Cooperation*. [conference paper]. IR 113. Vrije Universiteit, Amsterdam.
- Halfaker, A. and Taraborelli, D. (2015, November 30) Artificial intelligence service "ORES" gives Wikipedians X-ray specs to see through bad edits. *Diff*.
<https://diff.wikimedia.org/2015/11/30/artificial-intelligence-x-ray-specs/>
- Haraway, D. (1988). Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies*, 14(3), 575. <https://doi.org/10.2307/3178066>
- Harding, S. (1992). Rethinking standpoint epistemology – what is 'strong objectivity'? *The Centennial Review*, Vol. 36, No. 3 (FALL 1992), pp. 437-470.

- Hart, D. K. (1974). Social Equity, Justice, and the Equitable Administrator. *Public Administration Review*, 34(1), 3. <https://doi.org/10.2307/974395>
- Hartelius, E. J. (2010). *Wikipedia* and the Emergence of Dialogic Expertise. *Southern Communication Journal*, 75(5), 505–526. <https://doi.org/10.1080/10417940903377169>
- Hayles, K. (2002). Material Metaphors, Technotexts, and Media-Specific Analysis. In *Writing Machines*. MIT Press.
- Heidegger, M. (1977). *The question concerning technology, and other essays*. Garland Pub.
- Hemmendinger, D. (2024). HTML. Encyclopedia Britannica. <https://www.britannica.com/technology/HTML>
- Hill, B. M., & Shaw, A. (2013). The Wikipedia Gender Gap Revisited: Characterizing Survey Response Bias with Propensity Score Estimation. *PLoS ONE*, 8(6), e65782. <https://doi.org/10.1371/journal.pone.0065782>
- Hochschild, A. R. (2012). *The managed heart: Commercialization of human feeling*. University of California Press.
- Holmberg, L. (2021). Why Explore Stranded Encyclopedias? In L. Holmberg & M. Simonsen (Eds.), *Stranded Encyclopedias, 1700–2000* (pp. 3–34). Springer International Publishing. https://doi.org/10.1007/978-3-030-64300-3_1
- Ihde, D. (1990). *Technology and the Lifeworld: From Garden to Earth*. Indiana University Press.
- Ihde, D., & Malafouris, L. (2019). Homo faber Revisited: Postphenomenology and Material Engagement Theory. *Philosophy & Technology*, 32(2), 195–214. <https://doi.org/10.1007/s13347-018-0321-7>

- Iliadis, A., & Pedersen, I. (2018). The fabric of digital life: Uncovering sociotechnical tradeoffs in embodied computing through metadata. *Journal of Information, Communication and Ethics in Society*, 16(3), 311–327. <https://doi.org/10.1108/IICES-03-2018-0022>
- Ivory, C. (2013). The role of the imagined user in planning and design narratives. *Planning Theory*, 12(4), 425–441. <https://doi.org/10.1177/1473095212470546>
- Ivory, C., & Alderman, N. (2009). The imagined user in projects: Articulating competing discourses of space and knowledge work. *ephemera: theory & politics in organization*, 9(2), 131–148.
- Jackson, S. L. (1977). Towards a History of the Encyclopedia: From Amenemope of Egypt to the Collapse of Greek in Rome. *The Journal of Library History* (1974-1987), 12(4), 342–358.
- Jemielniak, D. (2014). *Common knowledge*. Stanford University Press.
- Jensen, M. M., & Aagaard, J. (2018). A postphenomenological method for HCI research. *Proceedings of the 30th Australian Conference on Computer-Human Interaction*, 242–251. <https://doi.org/10.1145/3292147.3292170>
- Jesper Simonsen, & Robertson, T. (2013). Participatory design: an introduction. In Jesper Simonsen, & Robertson, T. (Eds.), *Routledge International Handbook of Participatory Design*. Routledge.
- Joyce, R. A., & Tringham, R. E. (2007). Feminist Adventures in Hypertext. *Journal of Archaeological Method and Theory*, 14(3), 328–358. <https://doi.org/10.1007/s10816-007-9036-2>
- Kaur, B. (2012). Equity and social justice in teaching and teacher education. *Teaching and Teacher Education*, 28(4), 485–492. <https://doi.org/10.1016/j.tate.2012.01.012>

- Kendall, L. (1999a). Nerd nation: Images of nerds in US popular culture. *International Journal of Cultural Studies*, 2(2), 260–283. <https://doi.org/10.1177/136787799900200206>
- Kendall, L. (1999b). ‘The Nerd Within’: Mass Media and the Negotiation of Identity Among Computer-Using Men. *The Journal of Men’s Studies*, 3(7).
- Kennedy, H. (2012). *Net Work: Ethics and Values in Web Design*. Palgrave Macmillan.
- Kirkpatrick, G. (2003). Modernism and the Esthetics of Personal Computing. *Journal for Cultural Research*, 7(2), 179–194. <https://doi.org/10.1080/14797580305361>
- Konieczny, P. (2009). Governance, Organization, and Democracy on the Internet: The Iron Law and the Evolution of Wikipedia. *Sociological Forum*, 24(1), 162–192. <https://doi.org/10.1111/j.1573-7861.2008.01090.x>
- König, J., & Woolf, G. (2013). Introduction. In König, J., & Woolf, G. (Eds.), *Encyclopaedism from Antiquity to the Renaissance*. Cambridge University Press.
- Kuiper, K. (2024, February 19). Modernism. *Encyclopedia Britannica*. <https://www.britannica.com/art/Modernism-art>
- Kukowska, K. (2022). Knowledge Sharing Co-operative Structures: The Case of WikiProjects. *European Conference on Knowledge Management*, 23(1), 696–704. <https://doi.org/10.34190/eckm.23.1.570>
- Kuznetsov, S. (2006). Motivations of contributors to Wikipedia. *ACM SIGCAS Computers and Society*, 36(2), 1. <https://doi.org/10.1145/1215942.1215943>

- Lai, C.-Y., & Yang, H.-L. (2014). The reasons why people continue editing Wikipedia content – task value confirmation perspective. *Behaviour & Information Technology*, 33(12), 1371–1382.
<https://doi.org/10.1080/0144929X.2014.929744>
- Latour, B. (1990). Technology is Society Made Durable. *The Sociological Review*, 38(1_suppl), 103–131. <https://doi.org/10.1111/j.1467-954X.1990.tb03350.x>
- Latour, B. (2000). *Pandora's hope: Essays on the reality of science studies* (2. print). Harvard University Press.
- Latour, B. (2005). *Reassembling the social*. Oxford University Press.
- Leach, J., Nafus, D., & Krieger, B. (2009). Freedom Imagined: Morality and Aesthetics in Open Source Software Design. *Ethnos*, 74(1), 51–71.
<https://doi.org/10.1080/00141840902751188>
- Lemonnier, P. (Ed.). (1993). *Technological Choices: Transformation in material cultures since the Neolithic*. Routledge.
- Lessig, L. (2006). *Code* (Version 2.0). Basic Books.
- Levenson, M. (2011). *Modernism*. Yale University Press.
- Levy, S. (1994). *Hackers: Heroes of the computer revolution*. O'Reilly.
- Li, A. & Farzan, R. (2018). Keeping up on Current Events! A Case Study of Newcomers to Wikipedia. In Staab, S., Koltsova, O., & Ignatov, D. I. (Eds.), *Social Informatics: 10th International Conference, SocInfo 2018, St. Petersburg, Russia, September 25-28, 2018, Proceedings, Part I* (Vol. 11185). Springer International Publishing.
<https://doi.org/10.1007/978-3-030-01129-1>

- Lih, A. (2009). *The Wikipedia Revolution*. Aurum Press.
- Longo, F., Padovano, A., & Umbrello, S. (2020). Value-Oriented and Ethical Technology Engineering in Industry 5.0: A Human-Centric Perspective for the Design of the Factory of the Future. *Applied Sciences*, 10(12), 4182. <https://doi.org/10.3390/app10124182>
- Loos, A. (1908/2019). *Ornament and crime*. Penguin.
- Lundmark, S. (2018). Design project failures: Outcomes and gains of participation in design. *Design Studies*, 59, 77–94. <https://doi.org/10.1016/j.destud.2017.07.002>
- Marino, M. C. (2020). *Critical Code Studies*. MIT Press.
- Mateas, M., & Montfort, N. (2005). *A Box, Darkly: Obfuscation, Weird Languages, and Code Aesthetics*. [conference paper] Digital Arts and Culture: Digital Experience: Design, Aesthetics, Practice (DAC 2005). Copenhagen, Denmark, 2005.
- Mercier, H. (2016). The Argumentative Theory: Predictions and Empirical Evidence. *Trends in Cognitive Sciences*, 20(9), 689–700. <https://doi.org/10.1016/j.tics.2016.07.001>
- Mercier, H., & Sperber, D. (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences*, 34(2), 57–74. <https://doi.org/10.1017/S0140525X10000968>
- Merton, R. K. (1942/1973) The Normative Structure of Science. In Storer, N. W. (Ed.) *The Sociology of Science: Theoretical and Empirical Investigations*. The University of Chicago Press.
- Miller, J. K., Friedman, B., & Jancke, G. (2007). Value tensions in design: The value sensitive design, development, and appropriation of a corporation's groupware system. *Proceedings of the*

- 2007 *International ACM Conference on Conference on Supporting Group Work - GROUP*
'07, 281. <https://doi.org/10.1145/1316624.1316668>
- Moerenhout, T., Fischer, G. S., & Devisch, I. (2020). The elephant in the room: A postphenomenological view on the electronic health record and its impact on the clinical encounter. *Medicine, Health Care and Philosophy*, 23(2), 227–236.
<https://doi.org/10.1007/s11019-019-09923-5>
- Montfort, N., & Bogost, I. (2009). *Racing the Beam: The Atari Video Computer System*. MIT Press.
- Morgan, J. T., & Halfaker, A. (2018). Evaluating the impact of the Wikipedia Teahouse on newcomer socialization and retention. *Proceedings of the 14th International Symposium on Open Collaboration*, 1–7. <https://doi.org/10.1145/3233391.3233544>
- Nelson, T. H. (1974). *Dream machines*. Self-published.
- Niederer, S., & Van Dijck, J. (2010). Wisdom of the crowd or technicity of content? Wikipedia as a sociotechnical system. *New Media & Society*, 12(8), 1368–1387.
<https://doi.org/10.1177/1461444810365297>
- Norman, D. A. (2013). *The design of everyday things* (Revised and expanded edition). Basic Books.
- Novaes, C. D. (2015). A Dialogical, Multi-Agent Account of the Normativity of Logic. *Dialectica*, 69(4), 587–609. <https://doi.org/10.1111/1746-8361.12118>
- Novaes, C. D. (2018). The enduring enigma of reason. *Mind & Language*, 33(5), 513–524.
<https://doi.org/10.1111/mila.12174>
- Novaes, C. D. (2020a). The Role of Trust in Argumentation. *Informal Logic*, 40(2).
<https://doi.org/10.22329/il.v40i2.6328>

- Novaes, C. D. (2020b). Who's Afraid of Adversariality? Conflict and Cooperation in Argumentation. *Topoi*, 1–14. <https://doi.org/10.1007/s11245-020-09736-9>
- Novalis, & Wood, D. W. (2007). *Notes for a romantic encyclopaedia: Das Allgemeine Brouillon*. State University of New York Press.
- O'Sullivan, D. (2011), in
- Østergaard, K. L., Simonsen, J., & Karasti, H. (2018). Examining situated design practices: Nurses' transformations towards genuine participation. *Design Studies*, 59, 37–57. <https://doi.org/10.1016/j.destud.2017.12.002>
- Ottemo, A. (2019). The coproduction of masculinity and technology. In L. Gottzén, U. Mellström, & T. Shefer (Eds.), *Routledge International Handbook of Masculinity Studies* (1st ed., pp. 341–350). Routledge. <https://doi.org/10.4324/9781315165165-34>
- Pallardy, R. (2019). Encarta. Encyclopedia Britannica. <https://www.britannica.com/topic/Encarta>
- Papson, S., Goldman, R., & Kersey, N. (2004). Web Site Design: Hypertext Aesthetics and Visual Sociology. *American Behavioral Scientist*, 47(12), 1617–1643. <https://doi.org/10.1177/0002764204266241>
- Parsons, G. (2015). *The Philosophy of Design*. Cambridge: Polity books.
- Patton, A. (2003). Digital Aesthetics: The Legacy of Twentieth Century Modernism. *Technology Review*, 22 (Fall 2003).
- Peake, B. (2015). WP:THREATENING2MEN: Misogynist Infopolitics and the Hegemony of the Asshole Consensus on English Wikipedia. *Ada: A Journal of Gender, New Media, and Technology*, 4. <https://adanewmedia.org/2015/04/issue7-peake/>

- Peter, F. (2001). Health Equity and Social Justice. *Journal of Applied Philosophy*, Vol. 18, No. 2 (2001), pp. 159-170.
- Pinch, T. J., & Bijker, W. E. (1984). The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology might Benefit Each Other. *Social Studies of Science*, 14(3).
- Poulsen, K. S. (2019). Enactive individuation: Technics, temporality and affect in digital design and fabrication. *Phenomenology and the Cognitive Sciences*, 18(1), 1-18.
<https://doi.org/10.1007/s11097-017-9539-6>
- Rand, A. (1943). *The Fountainhead*. Macmillan.
- Rand, A. (1967/1990). *Introduction to Objectivist Epistemology*. New York: Penguin Books.
- Raymond, E. S. (2000). The Cathedral and the Bazaar. *First Monday*.
<https://firstmonday.org/ojs/index.php/fm/article/download/578/499?inline=1>
- Reagle, J. M. (2010). *Good Faith Collaboration*. MIT Press.
- Reinhart, K. M. (2022). A Cyclops, a Stone, and a Four-Breasted Woman: Drawing Anatomical Knowledge. *KNOW: A Journal on the Formation of Knowledge*, 6(1), 47-79.
<https://doi.org/10.1086/718566>
- Rosenberger, R. (2020). On hostile design: Theoretical and empirical prospects. *Urban Studies*, 57(4), 883-893. <https://doi.org/10.1177/0042098019853778>
- Rosenberger, R., & Verbeek, P.-P. (Eds.). (2015). *Postphenomenological investigations: Essays on human-technology relations*. Lexington Books.
- Russo, F. (2023). *Techno-Scientific Practices*. An Informational Approach. Rowman & Littlefield.

Sanger, L. (2005). *The Early History of Nupedia and Wikipedia: A Memoir—Slashdot*.

<https://features.slashdot.org/story/05/04/18/164213/the-early-history-of-nupedia-and-wikipedia-a-memoir>

Schleier, M. (2002). Ayn Rand and King Vidor's Film 'The Fountainhead': Architectural

Modernism, the Gendered Body, and Political Ideology. *Journal of the Society of Architectural Historians*, 61(3), 310–331. <https://doi.org/10.2307/991785>

Sennett, R. (2008). *The Craftsman*. Yale University Press.

<https://ebookcentral.proquest.com/lib/ucl/detail.action?docID=3420443>

Simon, B. (2007). Geek Chic: Machine Aesthetics, Digital Gaming, and the Cultural Politics of the

Case. *Mod. Games and Culture*, 2(3), 175–193.

<https://doi.org/10.1177/1555412007304423>

Simondon, Gilbert. (1966). *L'individu et sa g nese psycho-biologique*. J r me Millon.

Smits, M., Goor, H. van, Kallewaard, J.-W., Verbeek, P.-P., & Ludden, G. D. S. (2022). Evaluating

value mediation in patients with chronic low-back pain using virtual reality: Contributions

for empirical research in Value Sensitive Design. *Health and Technology*, 1–14.

<https://doi.org/10.1007/s12553-022-00671-w>

Steinmetz, K. F. (2015). Craft(y)ness: An Ethnographic Study of Hacking. *British Journal of*

Criminology, 55(1), 125–145. <https://doi.org/10.1093/bjc/azu061>

Stevenson, D. (1988). *The origins of freemasonry*. Cambridge University Press. [https://www-](https://www-cambridge-org.libproxy.ucl.ac.uk/core/books/origins-of-freemasonry/EB839D576657C1FBAD7BC26965468B30)

[cambridge-org.libproxy.ucl.ac.uk/core/books/origins-of-](https://www-cambridge-org.libproxy.ucl.ac.uk/core/books/origins-of-freemasonry/EB839D576657C1FBAD7BC26965468B30)

[freemasonry/EB839D576657C1FBAD7BC26965468B30](https://www-cambridge-org.libproxy.ucl.ac.uk/core/books/origins-of-freemasonry/EB839D576657C1FBAD7BC26965468B30)

- Sundin, O. (2011). Janitors of knowledge: Constructing knowledge in the everyday life of Wikipedia editors. *Journal of Documentation*, 67(5), 840–862.
<https://doi.org/10.1108/00220411111164709>
- Taraborelli, D., & Ciampaglia, G. L. (2010). Beyond notability. Collective deliberation on content inclusion in Wikipedia. 2010 *Fourth IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshop*, 122–125. <https://doi.org/10.1109/sasow.2010.26>
- Taylor, N., Hurley, U., & Connolly, P. (2016). Making Community: The Wider Role of Makerspaces in Public Life. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 1415–1425. <https://doi.org/10.1145/2858036.2858073>
- Thurow, L. C. (1981). Equity, Efficiency, Social Justice, and Redistribution. *Nebraska Journal of Economics and Business*, Spring, 1981, Vol. 20, No. 2 (Spring, 1981), pp. 5-24.
- Tkacz, N. (2011). The Politics of Forking Paths, in Lovink, G., & Tkacz, N. (Eds.), *Critical point of view: A Wikipedia reader*. Inst. of Network Cultures.
- Tkacz, N. (2015). *Wikipedia and the Politics of Openness*. The University of Chicago Press.
- Tsvetkova, M., García-Gavilanes, R., Floridi, L., & Yasseri, T. (2017). Even good bots fight: The case of Wikipedia. *PLOS ONE*, 12(2), e0171774.
<https://doi.org/10.1371/journal.pone.0171774>
- Turi, G. (1979). Ideologia e cultura del fascismo nello specchio dell'Enciclopedia Italiana. *Studi Storici*, 20(1), 157–211. <http://www.jstor.org/stable/20564602>
- Turkle, S. (1984). *The Second Self*. MIT Press.
- Turkle, S. (1995). *Life on the screen*. New York: Touchstone.

- Van de Kaa, G., Rezaei, J., Taebi, B., Van de Poel, I., & Kizhakenath, A. (2020). How to Weigh Values in Value Sensitive Design: A Best Worst Method Approach for the Case of Smart Metering. *Science and Engineering Ethics*, 26(1), 475–494.
<https://doi.org/10.1007/s11948-019-00105-3>
- Van De Poel, I. (2015). Design for Values in Engineering. In J. Van Den Hoven, P. E. Vermaas, & I. Van De Poel (Eds.), *Handbook of Ethics, Values, and Technological Design* (pp. 667–690). Springer Netherlands. https://doi.org/10.1007/978-94-007-6970-0_25
- Van De Poel, I. (2021). Design for value change. *Ethics and Information Technology*, 23(1), 27–31.
<https://doi.org/10.1007/s10676-018-9461-9>
- Van Dijck, J. (2013). *The Culture of Connectivity: A Critical History of Social Media*. Oxford Academic.
- Vee, A. (2017). *Coding literacy: How computer programming is changing writing*. The MIT Press.
- Verbeek, P.-P. (2005). *What Things Do Philosophical Reflections on Technology, Agency, and Design*. Penn State University Press.
- Verbeek, P.-P. (2011). *Moralizing Technology: Understanding and Designing the Morality of Things*. The University of Chicago Press.
- Wajcman, J., & MacKenzie, D. (Eds.). (1999). *The Social Shaping of Technology*. Open University Press.
- Wang, Y. (2023). Balancing Trustworthiness and Efficiency in Artificial Intelligence Systems: An Analysis of Tradeoffs and Strategies. *IEEE Internet Computing*, 27(6), 8–12.
<https://doi.org/10.1109/MIC.2023.3303031>

- Wertheim, M. (2000). *The Pearly Gates of Cyberspace: A History of Space from Dante to the Internet*. W.W. Norton.
- West, A. G., Kannan, S., & Lee, I. (2010). Detecting Wikipedia Vandalism via Spatio-Temporal Analysis of Revision Metadata. *EUROSEC '10: Proceedings of the Third European Workshop on System Security*.
- Wikimedia Foundation (2011). *Wikipedia editors study: Results from the editor survey, April 2011*.
- Wikimedia Foundation (2023). Sound logo: The Sound of All Human Knowledge
<https://wikimediafoundation.org/about/soundlogo/>
- Winner, L. (1980). Do Artifacts Have Politics? *Daedalus*, 109(1), 121–136.
- Witteman, H. O., Scherer, L. D., Gavaruzzi, T., Pieterse, A. H., Fuhrel-Forbis, A., Chipenda Dansokho, S., Exe, N., Kahn, V. C., Feldman-Stewart, D., Col, N. F., Turgeon, A. F., & Fagerlin, A. (2016). Design Features of Explicit Values Clarification Methods: A Systematic Review. *Medical Decision Making*, 36(4), 453–471.
<https://doi.org/10.1177/0272989X155626397>
- Wray, K. B. (2009). The Epistemic Cultures of Science and *Wikipedia*: A Comparison. *Episteme*, 6(1), 38–51. <https://doi.org/10.3366/E1742360008000531>
- Younkins, E. W. (2017). Atlas Shrugged and Social Change. *The Journal of Ayn Rand Studies*, 17(2), 285–305. <https://doi.org/10.5325/jaynrandstud.17.2.0285>
- Zhang, A. F., Wang, R., Blohm, E., Budak, C., Jr, L. P. R., & Romero, D. M. (2019). Participation of New Editors after Times of Shock on Wikipedia. *Proceedings of the International AAAI*

Conference on Web and Social Media, 13, 560–571.

<https://doi.org/10.1609/icwsm.v13i01.3253>

List of cited pages from Wikipedia and related wikis

Abstract Wikipedia [project page], https://meta.wikimedia.org/wiki/Abstract_Wikipedia

Conflict-of-interest editing on Wikipedia, https://en.wikipedia.org/wiki/Conflict-of-interest_editing_on_Wikipedia

Foundation: Legal: Visual identity guidelines,

https://foundation.wikimedia.org/wiki/Legal:Visual_identity_guidelines

Good faith, https://en.wikipedia.org/wiki/Good_faith

https://www.mediawiki.org/wiki/Wikimedia_Engineering_Architecture_Principles

Jimmy Wales, https://en.wikipedia.org/wiki/Jimmy_Wales

MediaWiki: Errors in the Wikipedia logo,

https://meta.wikimedia.org/wiki/Errors_in_the_Wikipedia_logo

MediaWiki: Experience [white paper] form 2019b

https://www.mediawiki.org/wiki/Wikimedia_Product/Perspectives/Experience

MediaWiki: Manual: MediaWiki architecture,

https://www.mediawiki.org/wiki/Manual:MediaWiki_architecture

MediaWiki: Phabricator/Help, <https://www.mediawiki.org/wiki/Phabricator/Help>

MediaWiki: Scale [white paper] form 2019a

https://www.mediawiki.org/wiki/Wikimedia_Product/Perspectives/Scale

MediaWiki: Tools [white paper],

https://www.mediawiki.org/wiki/Wikimedia_Product/Perspectives/Tools

MediaWiki: Wikimedia Engineering Architecture Principles,

https://www.mediawiki.org/wiki/Wikimedia_Engineering_Architecture_Principles

Meta: Anonymous users should not be allowed to edit articles,

https://meta.wikimedia.org/wiki/Anonymous_users_should_not_be_allowed_to_edit_articles

Meta: Association of Good Faith Wikipedians Who Remain Unregistered on Principle,

https://meta.wikimedia.org/wiki/Association_of_Good_Faith_Wikipedians_Who_Remain_Unregistered_on_Principle

Meta: Darwikinism, <https://meta.wikimedia.org/wiki/Darwikinism>

Meta: Special: Statistics, <https://meta.wikimedia.org/wiki/Special:Statistics>

Meta: Stewards, <https://meta.wikimedia.org/wiki/Stewards>

Meta: Unregistered user, https://meta.wikimedia.org/wiki/Unregistered_user

Meta: Wiki Loves Africa, https://meta.wikimedia.org/wiki/Wiki_Loves_Africa

MetaWiki: Logo/History, <https://meta.wikimedia.org/wiki/Logo/History>

MetaWiki: Movement Strategy, https://meta.wikimedia.org/wiki/Movement_Strategy

MetaWiki: Musings about unregistered contributors,

https://meta.wikimedia.org/wiki/Musings_about_unregistered_contributors

MetaWiki:Eventualism, <https://meta.wikimedia.org/wiki/Eventualism>

Special: Gadgets, <https://en.wikipedia.org/wiki/Special:Gadgets>

User:Gurch, <https://en.wikipedia.org/wiki/User:Gurch>

User:Jimbo_Wales/Statement_of_principles,

https://en.wikipedia.org/wiki/User:Jimbo_Wales/Statement_of_principles

User:Ram-Man, <https://en.wikipedia.org/wiki/User:Ram-Man>

Wikipedia: About, <https://en.wikipedia.org/wiki/Wikipedia:About>

Wikipedia: Blocking policy, https://en.wikipedia.org/wiki/Wikipedia:Blocking_policy

Wikipedia: Bot policy, https://en.wikipedia.org/wiki/Wikipedia:Bot_policy

Wikipedia: Civility, <https://en.wikipedia.org/wiki/Wikipedia:Civility>

Wikipedia: Conflict of interest, https://en.wikipedia.org/wiki/Wikipedia:Conflict_of_interest

Wikipedia: Copyrights, <https://en.wikipedia.org/wiki/Wikipedia:Copyrights>

Wikipedia: Disruptive editing, https://en.wikipedia.org/wiki/Wikipedia:Disruptive_editing

Wikipedia: Don't bite the newcomers,

https://en.wikipedia.org/wiki/Wikipedia:Please_do_not_bite_the_newcomers

Wikipedia: Don't bite the researchers,

https://en.wikipedia.org/wiki/Wikipedia:Don%27t_bite_the_researchers

Wikipedia: Editing on mobile devices,

https://en.wikipedia.org/wiki/Wikipedia:Editing_on_mobile_devices

Wikipedia: Editing policy, https://en.wikipedia.org/wiki/Wikipedia:Editing_policy

Wikipedia: Editors should be logged-in users,

[https://en.wikipedia.org/wiki/Wikipedia:Editors_should_be_logged-in_users_\(failed_proposal\)](https://en.wikipedia.org/wiki/Wikipedia:Editors_should_be_logged-in_users_(failed_proposal))

Wikipedia: Esperanza, <https://en.wikipedia.org/wiki/Wikipedia:Esperanza>

Wikipedia: Ethically researching Wikipedia,

https://en.wikipedia.org/wiki/Wikipedia:Ethically_researching_Wikipedia

Wikipedia: Etiquette , <https://en.wikipedia.org/wiki/Wikipedia:Etiquette>

Wikipedia: Five Pillars, https://en.wikipedia.org/wiki/Wikipedia:Five_pillars

Wikipedia: Gaming the System, https://en.wikipedia.org/wiki/Wikipedia:Gaming_the_system

Wikipedia: History of Wikipedia bots,

https://en.wikipedia.org/wiki/Wikipedia:History_of_Wikipedia_bots

Wikipedia: History of Wikipedia, https://en.wikipedia.org/wiki/History_of_Wikipedia

Wikipedia: Huggle, <https://en.wikipedia.org/wiki/Wikipedia:Huggle>

Wikipedia: Ignore all rules, https://en.wikipedia.org/wiki/Wikipedia:Ignore_all_rules

Wikipedia: Image use policy, https://en.wikipedia.org/wiki/Wikipedia:Image_use_policy

Wikipedia: IP addresses are not people,

https://en.wikipedia.org/wiki/Wikipedia:IP_addresses_are_not_people

Wikipedia: IPs are human too,

https://en.wikipedia.org/wiki/Wikipedia:IP_editors_are_human_too

Wikipedia: List of Wikipedians by number of edits,

https://en.wikipedia.org/wiki/Wikipedia:List_of_Wikipedians_by_number_of_edits

Wikipedia: Magnus Manske Day,

https://en.wikipedia.org/wiki/Wikipedia:Magnus_Manske_Day

Wikipedia: Manual of Style, https://en.wikipedia.org/wiki/Wikipedia:Manual_of_Style

Wikipedia: Mobile communication bugs,

https://en.wikipedia.org/wiki/Wikipedia:Mobile_communication_bugs

Wikipedia: Neutral point of view,

https://en.wikipedia.org/wiki/Wikipedia:Neutral_point_of_view

Wikipedia: No original research, https://en.wikipedia.org/wiki/Wikipedia:No_original_research

Wikipedia: Oral citations experiment,

https://en.wikipedia.org/wiki/Wikipedia:Oral_citations_experiment

Wikipedia: Policies and guidelines,

https://en.wikipedia.org/wiki/Wikipedia:Policies_and_guidelines

Wikipedia: Protection policy, https://en.wikipedia.org/wiki/Wikipedia:Protection_policy

Wikipedia: Requests for adminship,

https://en.wikipedia.org/wiki/Wikipedia:Requests_for_adminship

Wikipedia: Research, <https://en.wikipedia.org/wiki/Wikipedia:Research>

Wikipedia: STiki, <https://en.wikipedia.org/wiki/Wikipedia:STiki>

Wikipedia: Templates, <https://en.wikipedia.org/wiki/Wikipedia:Templates>

Wikipedia: Types of bots, https://en.wikipedia.org/wiki/Wikipedia:Types_of_bots

Wikipedia: User access levels, https://en.wikipedia.org/wiki/Wikipedia:User_access_levels

Wikipedia: Vandalism, <https://en.wikipedia.org/wiki/Wikipedia:Vandalism>

Wikipedia: Verifiability, <https://en.wikipedia.org/wiki/Wikipedia:Verifiability>

Wikipedia: Verifiability, <https://en.wikipedia.org/wiki/Wikipedia:Verifiability>

Wikipedia: Village Pump (idea lab)/Archive 33,

[https://en.wikipedia.org/wiki/Wikipedia:Village_pump_\(idea_lab\)/Archive_33](https://en.wikipedia.org/wiki/Wikipedia:Village_pump_(idea_lab)/Archive_33)

Wikipedia: VisualEditor, <https://en.wikipedia.org/wiki/Wikipedia:VisualEditor>

Wikipedia: Welcome unregistered editing,

https://en.wikipedia.org/wiki/Wikipedia:Welcome_unregistered_editing

Wikipedia: What are these researchers doing in my Wikipedia?,

https://en.wikipedia.org/wiki/Wikipedia:What_are_these_researchers_doing_in_my_Wikipedia%3F

Wikipedia: What Wikipedia is not,

https://en.wikipedia.org/wiki/Wikipedia:What_Wikipedia_is_not

Wikipedia: Who writes Wikipedia?,

https://en.wikipedia.org/wiki/Wikipedia:Who_writes_Wikipedia%3F

Wikipedia: Wikipedia is a work in progress,

https://en.wikipedia.org/wiki/Wikipedia:Wikipedia_is_a_work_in_progress#:~:text=This%20page%20is%20a%20nutshell,a%20constant%20work%20is%20progress.

Wikipedia: Wikipedia logos, https://en.wikipedia.org/wiki/Wikipedia:Wikipedia_logos

Wikipedia: WikiProject Women in Red,

https://en.wikipedia.org/wiki/Wikipedia:WikiProject_Women_in_Red

Wikipedia: Your first article, https://en.wikipedia.org/wiki/Help:Your_first_article

WikiTrust, <https://en.wikipedia.org/wiki/WikiTrust>

Appendix A: Epistemic culture themes

Theme 1 – Neutrality as objectivity

The central epistemic principle animating Wikipedia is neutrality, i.e. a form of objectivity which is conceived as not taking a side on debates about the truth of the matter, using neutral language, and ignoring positionality. Neutrality as a guiding principle is apparent across contexts, assuming a different form depending on context. Policy defines neutrality as a central tenet of Wikipedia's epistemology, defined as Neutral Point of View. Neutrality as an epistemic guide is also often evoked by Wikipedians themselves, in conversations on talk pages, in interviews, and in informal writings such as essays and guidelines. In conversation, both among themselves and with me, Wikipedians often display concern for "bias", understood as a deviation from a neutral point of view. Affective neutrality – invitations to keep conversations low-affect, not edit while one is in a bad mood, not using emotionally charged language – is also both prescribed by policy and practiced by Wikipedians themselves.

Theme 2 – Newcomers and Wikipedia's preservation

Similarly to neutrality, the twin preoccupations of expanding participation and guaranteeing the survival of Wikipedia are present across contexts. Wikipedia's design is reliant on large numbers of editors working to produce and update content: since anyone can edit Wikipedia, expertise can't be counted on in order to gather and maintain content. Since non-experts are prone to error, it is important that mistakes are spotted quickly: for Wikipedia to be self-correcting, many people need to be looking at its content at the same time. At the same time, because Wikipedia is run by volunteers, and content grows fast, there needs to be an ever-increasing group of people willing to work on the project, able to keep up with Wikipedia's growth, and replace those who are no longer able or willing to participate. Treating newcomers with grace is deemed necessary for the survival of the project.

Theme 3 – Individuality and community

Wikipedia acts both as a crowd – a collection of unconnected individuals – and a community. As a whole, those who edit Wikipedia are not necessarily connected to each other in lasting relationships; rather, the Wikipedia process is designed exactly, as described above, to be potentially open to anyone. In principle, Wikipedia is a crowd, and that is consistent with the ideal of open participation: communities tend to protect themselves from outsiders, which Wikipedia is ostensibly trying to avoid. Policy and guidelines describe Wikipedia as a crowd. At the same time, large parts of Wikipedia function, and represent themselves, as a community. People get together to work on specific projects, talk to each other, form relationships, give each other awards for work on the encyclopaedia, and even for being a good co-editor. Individuality is also highlighted and celebrated, as exemplified by the "be bold" imperative, which invites Wikipedians to do what they think is right, without checking with others; the radically decentralised nature of Wikipedia means individual choice is the primary mover of action, if within a collective context.

Theme 4 – Brute force and negotiation

The process of editing Wikipedia is based on negative feedback and competition: editors are free to edit pages, on the basis that any error can be corrected by others (this is known as "radical collaboration"). In principle, as described in policies and guidelines, no contribution is more valuable or precious than others, and each edit should be evaluated in its own terms. The practice of editing Wikipedia, however, is heavily reliant on negotiation: Wikipedians debate edits in talk pages, form alliances over specific topics, call on each other for expertise and support. Wikipedians also implement strategies (known within the community as "gaming the system") in order to win arguments. Knowing how Wikipedia works, and how to use the mechanisms that regulate edits to one's advantage, is part and parcel of being an effective Wikipedia editor.

Theme 5 – Anonymity and identity

At this point in time, people who decide to edit Wikipedia can share as much personal information as they wish, ranging from anonymity, where the editor is known by their IP address, to their own real name, interests, real-life expertise. Some Wikipedians have a public profile related to their activity on Wikipedia: Jesse Wade, for instance, is known for her work on the coverage of women; Magnus Manske for coding work. Wikipedia's core ethos of wide participation, which doesn't require personal credentials, encourages anonymity – conflict often arises when personal experience or expertise are brought to bear during deliberation about the content of articles.

Discussion: The clean channel ideal

Looking at all the themes at a glance, a picture of the ideal knower emerges, which I have termed "the clean channel ideal". The ideal knower, in Wikipedia's culture, is capable of holding a neutral point of view, carrying information from secondary sources to the encyclopaedia without interfering, and finally maintain an attitude of distance from the content they produce. Consequently, the ideal knower ought to refrain from bringing their own positionality or expertise to bear when editing, take sides in debates over the truth of specific claims, and becoming emotionally attached to their own work. Disinterestedness and detachment, then, constitute the basic epistemic virtues valued on Wikipedia. It is interesting that these virtues imply a bracketing of the body, as the physical anchoring in a point of view; positionality – having a point of view, shaped by one's location in the world – is, as much as possible, ignored. Positionality carries the inclination to defend one's own community or interests – or, as Wikipedians like to call this, bias. Interests result, in part, from needs, such as the need for sustenance or belonging, which are supposed to be neglected by the ideal editor: volunteering, not receiving payment for services, implies ignoring the need for sustenance; editing autonomously means working without necessarily engaging with a community. Again, ignoring needs means ignoring the body. Finally, detachment, in an emotional sense, amounts to downplaying affect, which is also typically associated with bodiliness (while rationality, albeit incorrectly, tends to be conceptualised as separate from the body).

The clean channel ideal is embedded in a certain strain of Wikipedia's culture: with reference to the themes I generated from my analysis, it runs through the idea of objectivity as neutrality (Theme 1), and one end of the tensions I outlined through Themes 3, 4 and 5. Conceptualising Wikipedia as a crowd (Theme 3) means upholding disinterestedness, understood as absence of belonging. Radical

collaboration, necessary to brute force the encyclopaedia (Theme 4), implies emotional detachment. Finally, anonymity (Theme 5) speaks for itself: the literal hiding of personal feature naturally speaks to disembodiedness.

At the same time, the other ends of the tensions in Themes 3, 4, and 5 highlight a resistance to the clean channel ideal, in the forms of introducing, and tending to, relationships between Wikipedia community members, allowing for negotiation, which implies a more personal management of the editing process, and making space for the personalities of Wikipedians to be affirmed. The reason why I frame the clean channel ideal as the core of Wikipedia's epistemology, and these pockets of resistance as resistance is ideological and historical. Ideologically, the neutrality principle, as Wikipedia stands, doesn't have opponents: it is upheld enthusiastically by the community, and written in policy; Wikipedia leadership, especially the only founder who is still active Jimmy Wales, defend it with conviction; the policy cannot be easily changed (the page is protected, meaning it is very difficult to change the wording of the policy itself). Since detachment and disinterestedness are, essentially, corollaries of the Neutral Point of View policy, in my view the clean channel model is the core ideological tenet of Wikipedia's epistemology. I doubt most established members of the Wikipedia community would contest this point.

From a historical point of view, opposition to expressions of the clean channel ideal have appeared over time: the basic design for Wikipedia simply consisted in editing infrastructure. Conversations among Wikipedians were, by design, hidden from the main view. Wikipedians registered with a pseudonym, put very little information on their userpages, and kept interaction to a minimum: practices such as the introduction of Userboxes and the ability to give awards appeared later, and were initially controversial. Conversations around positionality and its meaning are recent and nowhere near settled.

Appendix B: Infrastructure and design themes

The themes below are discussed extensively in the body of the thesis, where I have also placed illustrative quotes and visuals. Here I provide a brief overview, for reference, to connect my analysis with empirical results.

Theme 1 – User experience

The experience of the reader and/or editor is central to the Wikipedian discourse on design. Specifically, the focus is on affect, aesthetics, and rewards. Affect is discussed in terms of what it feels like read or edit Wikipedia. As far as reading is concerned, the main considerations are around engagement – a key strategic aim of the Wikimedia Foundation is for Wikipedia to stay relevant and attract new readership; often, this priority translates in discussions of interest or curiosity. In terms of editing, affect is relevant in terms of the atmosphere on talk pages, and the effect it can have on retention of editors, as well as the reward system, discussed in terms of whether it provides sufficient positive feedback to editors. The aesthetics of the editing interface are relevant to experience in terms of their usability and accessibility.

Theme 2 – Ideals of modernity

Wikipedia has always, on one hand, chased an ideal of modernity, and on the other, hosted strains of conservatism and nostalgia. The worry around being obsolete is often present, paired with discussions, plans, imagined alternatives to the existing infrastructure that may be developed in the future. Said alternatives consist, on one hand, of multimedia content, under the assumption that a wider range of formats would appeal to a wider audience, or, in some cases, of reducing the dependence between Wikipedia's content and its underlying platform.

Theme 3 – Belonging

Underlying the concerns described above is the question of who Wikipedia should appeal to, expressed often by imagining users, either implicitly or explicitly. Implicitly, the user is imagined through design choices and discussions around choices, where the rationale is anchored to the envisioned skills or personal qualities of users. Explicit framings of users occur through offering access to users with specific skills and policing behaviour.

Theme 4 – Justice

Considerations around justice inform design choices along three main lines: accessibility, freedom, and ubiquity. Accessibility is framed mainly in relation to disability, as the quality of a given artefact that grants the ability to use the artefact to as wide a range of users as possible. Freedom is articulated in various versions, as self-determination – for instance, in the form of a right to customisation – free speech, and financial independence. Ubiquity concerns the dissemination of

Wikipedia's content; the ambition of the project is for information to be available to as many people as possible, beyond technological barriers such as availability of devices and Internet connection.

Appendix C: Codebook

I used inductive coding, guided by my research questions. As questions evolved over the course of my study, so did coding schemes. I present here a unified coding scheme, in its final form. I have grouped codes by area.

Group 1 – Aesthetics of content

Code	Definition	Example
Clarity	Appeals to think about the reader: importance of clarity and good communication.	I do agree there is room for making the last part clearer, since as it is currently written, it could be interpreted the same way as e.g. "non-binary people are those who are not men or women", i.e. as saying people have to be not (the thing on the left side of the or) and also not (the thing on the right side of the or). Perhaps it could be revised to "...who are not heterosexual or are not cisgender"?
Communication	This code applies to cases when communication is discussed as a value, both social and epistemic (the boundary between the two being rather blurry in this context).	The new name is an improvement over the old one, because it's informative and doesn't in itself imply anything negative about the main APIs and various avenues of access. It's also a term that companies are more likely to search for.
House style	Word choice, style, grammar.	Avoid stating opinions as facts. Usually, articles will contain information about the significant opinions that have been expressed about their subjects. However, these opinions should not be stated in Wikipedia's voice. Rather, they should be attributed in the text to particular sources, or where justified, described as widespread views, etc. For example, an article should not state that "genocide is an evil action", but it may state that "genocide has been described by John So-and-so as the epitome of human evil."

Languages, translation & linguistic proficiency	This code labels treatments of linguistic matters. These include discussions of the different languages in which Wikipedia is written, languages spoken by Wikipedians, English-language proficiency, what kind of spelling should be used, translation.	<p>I promise that the documentation pages for "Wikimedia Enterprise" (the Essay, FAQ, Principles, and Main page) will all be translated and published in Russian, next week.</p> <p>The information was all published on Tuesday in English, French, German, Spanish, and some Italian</p>
Neutral wording	Talking about wording in terms of neutrality or bias. Discussions or references to loaded language.	In some cases, the choice of name used for a topic can give an appearance of bias. While neutral terms are generally preferable, this must be balanced against clarity.
Visuals	This code applies to discussions of images included into Wikipedia, as illustrations of articles. Standards for quality of images, ways of choosing images.	When making user-made diagrams or similar images, try not to use color alone to convey information, as it is inaccessible in many situations.

Group 2 – Authorship & identity

Code	Definition	Example
Authorship/POV	Instances in which authorship is discussed, as well as the point of view that an article (or part of an article) appears to be written from, and the connection between an editor's identity and the content they produce.	Wikipedia's visibility makes it a natural haunt of viewpoint pushers on political and social controversies. Even if you try to be scrupulously careful about avoiding POV edits, other editors working on the same topic may assume that you are a party to the dispute and assign you to one of the various camps. If this offends, annoys, or troubles you, you should consider staying out of the fray.
Behaviour / etiquette	Prescriptions about how agents (editors, bots, other members of the community)	I have not pretended to be neutral on this subject. I have stated a position and am highly offended by much of the

	should interact with each other; what constitutes polite or rude behaviour.	discussion at the MfD. There is nothing improper about that.
Experience	Cases in which people mention their own experience and/or experience of others. Discussion of whether this should be allowed.	I agree with this point of view. Currently it would seem to me at least that Sedna has a greater claim on planethood than Pluto, given the uniqueness of its orbital location. I particularly agree with your last point there on classifying Pluto with Minor Planet number 0. I actually suggested that to Brian Marsden of the Minor Planet Centre about 7 years ago when the whole debate about classifying it as 10000 was floating about, and he indicated he had a favourable view of that solution.
Expertise	Signalling expertise, either explicitly ("I am a physicist") or implicitly ("you need to learn more algebra in order to understand that"). Discussion of this attitude. Importance/unimportance of expertise. Role of expertise on Wikipedia. What is expertise on Wikipedia.	EDIT: I am very sorry I did not read all of the articles, just skimmed their titles i am not an advanced astronomer(no astronomy classes sinse the 5th grade) so i am not familiar with hubble time or all of these big words, so let me get this straight, neptunes gravity is what make pluto cross its orbit so there fore pluto doesn't clear the neighborhood?
Personal identity	Cases in which people explicitly mention / gesture towards their IRL identity. Discussion over whether one's identity matters, technologies that allow expression of identity (see userboxes, "fun" usernames, etc.).	Can't we just get rid of all political userboxes? This is a giant timesink that contributes nothing to the project. How is it helpful to find out someone you enjoy collaborating with has political views you find abhorrent? Actually, let's just get rid of any userbox that isn't project-related. I don't need to know you went to UMich, either. :)

Group 3 – Community & policy

Code	Definition	Example
Aggressiveness / masculinity	Discussion and policing of displays of aggression and discussions of masculinity.	It does make sense that the hypernegative aspect of Wikipedia improves the odds that someone will take my serious questions as irrelevant insults, but it should also highlight the difference between my conversations starting as inquiries versus others' starting as condemnations.
Anonymous editors / IP edits	Treatment of those who edit anonymously (also known as 'IPs', as the IP address of those who edit anonymously is displayed alongside their edits and talk page posts). The way this category of editors is discussed, and discriminated against, elucidates the boundaries of the Wikipedian self as a knower, pressing on questions of identity and belonging.	Anonymous personage you really know nothing of me. Although my college degree is in history it's not like I'm some uneducated fool on this subject. In fact my history thesis concerns science to a certain degree and the history of science is an interest in mind. I can name several astronomy articles I created here if need be. However the analogies you are making are basically a form of Scientism. Things are not scientific merely because we like scientists, trust their judgement, and can find a justification for them. If they were then we'd still have to accept eugenics among other things. The whole notion is practically a debasement of science. Science is a proving or disproving of hypotheses. People are making out like rejecting this verdict is the same as rejecting evolution or the Big Bang. That idea is ridiculous and highly insulting. If you can't see that it's your problem. I'll stop there as a storm is coming in my region.
Argumentation	Discussion of argumentation, defined as the activity of debating with others, and its role within Wikipedia's knowledge-production process.	Many of these discussions will involve polls of one sort or another; but as consensus is determined by the quality of arguments (not by a simple counted majority), polls should be regarded as

		structured discussions rather than voting. Responses indicating individual explanations of positions using Wikipedia policies and guidelines are given the highest weight.
Conflict, disagreement	Instances where disagreement and conflict are discussed. This is related to argumentation, but while argumentation has a softer connotation, and is verbal in nature, this code tracks conversations about conflicts that may manifests in ways other than debate, such as hostile behaviour, inclusion of material aimed to offend, and so forth.	if the userboxes are "inflammatory and divisive", having a heated discussion like what we had, and then having such a discussion closed either way (be it "delete", "keep" or "no consensus") would clearly be much more "inflammatory and divisive" to the community, as both sides have strong arguments and more importantly, strongly held opinions.
Consensus	Role played by consensus, within Wikipedia: what it does, how it's defined.	Further analysis shows that consensus is often hard to reach. The figure below (left) shows the distribution of duration (in seconds) of AfD discussions. While the vast majority lasts less than one week, there is a sizable minority that takes up to two weeks or more. Some discussions took more than one year to be closed!
Disruption, trolling, vandalism	Discussion of practices that hamper Wikipedia's normal and desirable functioning.	My user page was vandalized a few dozen times (maybe only 15), I believe always homophobic, until I removed any picture of myself (presumably triggering antigay bigots' ire FAR less).
Internal hierarchy & status	Cases in which people appeal to their role/status within Wikipedia, or mention someone else's. Discussion of using Wikistatus. Talking about internal hierarchy and power within the community.	I remind everyone to be civil and assume good faith. Some of the comments in the MfD are so intemperate and outside the bounds of civil discourse that if I were not WP:INVOLVED I would have handed down formal warnings.
Internal policy	Relatively broad code, outlining the general policy infrastructure on Wikipedia,	What is settled is that per our existing policies, user pages are not platforms to express support for

	beyond encoded policy. Through conversations, discussion of policy status and how it should be determined, these data points fix the structure and boundaries of Wikipedia's internal policy.	discrimination. We are entirely capable of determining, on a case by case basis, whether a userbox falls afoul of WP:UBCR. Where it gets political is when people present say that we cannot evaluate that case because it is a "side". The "sides" are about politics, and are independent of our determination of whether or not something is discriminatory.
Number of editors / new editors	Worries around dwindling numbers of editors, recruitment of new editors. Claims that a specific feature will discourage newcomers.	A bottom line: don't assume nefarious intentions (including bias) when an article's shortcomings can be attributed to lack of resources/editors, lack of skill, or cultural/legal differences.
Politics	Mentions of politics and recent history.	If they define themselves by bigotry, sufficiently strongly to add a userbox, then absolutely. They go to Conservapedia. What would you do if anyone added a userbox saying that they believe marriage means two people of the same colour? This is exactly the same. And yes, I do mean exactly the same.
Procedure / bureaucracy	Detailed discussions of procedure. This code is mainly aimed at evidencing the level of attention dedicated by Wikipedians to bureaucracy.	Given that the idea that pages can change is at the very heart of the **wiki** movement, it seems that a consultation procedure that assumes stable text is misguided. I'd suggest updating the procedures to match reality, rather than expecting people to stop editing pages. That's especially true now that the one of the recommendations of the strategy process is that we manage internal knowledge - adding relevant information to pages is something that needs to be encouraged. I think it makes sense to link to as much on-wiki discussion as possible.
Safety / harassment	Discussion of what constitutes harassment within Wikipedia, how it should be handled, how it justifies	Not experiencing heavy harassment but knowing of situations where moderators did tend to act against complainers or play down complaints,

	behaviours and policy.	from the POV there are no local rules. Too often people simply leave the project after such an experience. Clear universal rules could be of help to broaden editing communities, therewith diversifying content and attract broader reader groups. I
Wikimedia Foundation	Direct mentions of the Wikimedia Foundation.	I can't remember when we last had a change to the rules about signatures, other than the WMF reform to them which has proved remarkably uncontentious for a WMF change.

Group 4 – Knowledge-production & epistemology

Code	Definition	Example
Dissemination of knowledge	Instances in which the dissemination of knowledge is invoked or used to support arguments. Communicating knowledge is Wikipedia's main aim, according to Wikipedia itself.	In the Movement Strategy recommendations to increase the sustainability of our movement and improve user experience there are the recommendations to, respectively: "Explore new opportunities for both revenue generation and free knowledge dissemination through partnerships and earned income—for example...Building enterprise-level APIs,
Feedback/brute force/evolutionary mechanisms	Often, accounts of knowledge-production on Wikipedia hinge on some kind of feedback mechanism. This code tracks similar views on Wikipedia's process, within Wikipedia itself.	Wikipedia is a demonstration of the classical liberal philosophy of Human Action, whereby optimized wealth emerges out of a decentralized, free and competitive process of trial-and-error creation.
Informal epistemology	Wikipedians reflect deeply on their practice. They hold beliefs and theories about the way Wikipedia works, and how it should work.	"There is no deadline."
Informal epistemology: Deletionism	Deletionism is an approach to editing Wikipedia. I have subsumed it under 'informal	Deletionism is a term used to describe the minimalist idea that removal of material is a productive mode

	epistemologies' – ways in which Wikipedians conceptualise their practice. I used this code to mark text about deletionism: explanations of what it is, opinions about it, articulations of it.	of editing, and that doing such regularly, primarily, or solely, rather than adding material, is sufficient to call oneself an editor. The main claim is that the wiki mode of editing is 'too quick', and that there are too many of them (contributors, of some nominal spectrum of skill), and that culling is therefore an art.
Informal epistemology: Inclusionism vs exclusionism	Inclusionism and exclusionism are opposing approaches to editing Wikipedia. Quotes labelled with this code elucidate what inclusionism and exclusionism are, how they relate to each other, what they entail in terms of editing practice.	There are no deletionists. There are inclusionists. That's because inclusionism is a somewhat sensible philosophy. It's somewhat sensible to think we should have an article on every possible topic, regardless of notability. There are no deletionists, though, because there is nobody who thinks that we should delete all articles. There's not even anybody who believes that we should delete all non-perfect articles. Hell, there's not even anybody that believes we should delete all stubs.
Informal epistemology: Wikipedia is not finished / eventualism	Connected with Wikipedia's informal epistemology, this code applies to uses of the idea that Wikipedia is not (and/or never will be) finished.	Eventualism is a tendency amongst Wikipedians focusing on the eventual value of Wikipedia in the long-term rather than the immediate value.
Knowledge-production process	Reflection and information about the process by which knowledge is produced on Wikipedia.	Many of these discussions will involve polls of one sort or another; but as consensus is determined by the quality of arguments (not by a simple counted majority), polls should be regarded as structured discussions rather than voting. Responses indicating individual explanations of positions using Wikipedia policies and guidelines are given the highest weight.
NOR	NOR means "No original research" and is one of Wikipedia's core policies. This code applies to text about the	In fact, that's not original research, and you can see that in other discussions, even references, but I know there's a strong anti-Pluto bias in wikipedia and

	policy, explaining the policy, or depling the policy in interaction.	views over categorization of celestial bodies, which does not reflect many important references.
Notability	Notability is one of the main criteria for inclusion of content on Wikipedia. This code mainly tracks how it is used in conversation, and defined.	To counter Wikipedia's systemic bias against covering women's biographies, women's work and women's perspectives, which occurs because the vast majority of editors are male (see the article on Gender bias in Wikipedia for more details), we have a department, Women in Red, that works on missing articles about notable women.
NPOV	Neutral Point of View is one of Wikipedia's core policies. This code labels instances in which the policy is defined, explicitly mentioned, or strongly implied.	@Blueboar: the problem is repeatedly that what appears to be a neutral tone to one editor/reader does not to another. I think that the articles on Christian Science and Traditional Chinese medicine are neutral – too neutral in places in my view – but I suspect that believers in either would take a different view. Certainly it has been a constant battle to get any criticism into some of the articles on traditional Indian medical systems, and believers are constantly adding bits to articles about plants used in these systems that violate WP:MEDRS and have to be removed. So the key question is "who defines neutral"?
Sources	Any mention of what a source is, what makes a source reliable or unreliable.	Citations to non-English reliable sources are allowed on the English Wikipedia. However, because this project is in English, English-language sources are preferred over non-English ones when available and of equal quality and relevance.
Strategy / gaming the system	Wikipedians often use rules and platform affordances to their advantage, in order to win arguments or influence the content of a page. Explicit discussion of these strategies fall	Bold editors of policy and guideline pages are strongly encouraged to follow WP:1RR or WP:0RR standards. Editing a policy to support your own argument in an active discussion may be seen as gaming the system, especially if

	under this code.	you do not disclose your involvement in the argument when making the edits.
Truth	Ways in which the concept of truth is discussed or deployed on Wikipedia.	Even if you are sure something is true, it must be verifiable before you can add it.
Verifiability	Verifiability is one of Wikipedia's core policies. This code applies when V is discussed, defined, explicitly mentioned or strongly implied.	Do you have any sources for this?

Group 5 – Subjectivity & the knower

Code	Definition	Example
Affect	Talking about feelings, emotions, affective atmosphere. Expressing strong feelings or emotions. What feelings or emotions are encouraged/allowed.	Some of the comments posted are clearly inconsistent with the project's longstanding neutrality on divisive social/political issues. Suggesting that fellow Wikipedians who are opposed to SSM hold views equivalent to supporting slavery is shockingly offensive and a gross violation of WP:NPA and WP:CIVIL.
Civility	Civility is the only behavioural core policy. This code marks instances when the policy is defined and/or deployed in conversation.	I'm aware of that essay, and I've even made a minor contribution to it. I'm mainly wondering why the editors are sometimes impolite regarding pseudoscience.
Epistemic skills & virtues	Discussion of what skills are required in order to edit Wikipedia, what makes for a good editor.	Being able to use a markup language doesn't make you a good encyclopedia contributor. Being knowledgeable about what you're writing is key.
Technologies of the self	Practices aimed at building a good editor, including telling people what they should do. (Paired with Affect: telling people what they should feel.)	I remind everyone to be civil and assume good faith. Some of the comments in the MfD are so intemperate and outside the bounds of civil discourse that if I were not WP:INVOLVED I would have handed down formal warnings.
Wiki-virtues	Discussion of what kind of virtues should be developed by a	"She epitomized the comment Jimbo once made about how it amazed him that

	good editor, or what makes a good editor. Compliments. Praise.	WP editors could put up with so much and for as long as we do."
--	--	---

Group 6

Technology & design

Code	Definition	Example
Accessibility	Instances in which issues around accessibility and disability are brought up, in a technical context. While this code is related to "Diversity / equity / non-Western knowledges / epistemic injustice / inclusion", I preferred to keep it separate, because issues around accessibility tend to be framed as specifically technical (while other kinds of "diversity"-related considerations tend to be framed as social, or sociotechnical). Accessibility is special because it has a longer history, in terms of Web design ethics, than other, bordering, ethical concerns. At the beginning of the 00s, when Wikipedia was founded, accessibility was the main ethical concern in Web design (see: Net Work).	JPEG files should not use arithmetic coding due to limited browser support. Please prefer Huffman coding for JPEG files instead.
Customisation	Conversations around the importance of customisation. Mentioning customisation as a rationale for design choices.	Generally speaking and with exceptions, individual freedom trumps the greater good at en-wiki. Custom signatures are another example, where the community consistently refuses to impose even minimal bright-line restrictions.
Decentralisation	Framings of Wikipedia as a decentralised system.	A third powerful metaphor for Wikipedia is that it is actually an evolving ecosystem of ideas where only the best and strongest survive. This view extends the classical liberal

		philosophy of Human Action, whereby optimized value emerges out of a decentralized, free and competitive process of trial-and-error creation.
Editing tools & mechanisms	Explanations/discussions of editing infrastructure. How it works, how editors should or shouldn't use it, how it is used.	Edit summaries are useful tools when building an encyclopedia, and having them can hours and hours of work from the many editors that might have to review each edit to tell what it is.
Format	Mentions of the format in which the content of Wikipedia should be displayed. Discussion over which format is best. Strategic planning around format.	This new Wikimedia project will build a library of functions, written by volunteers, to help answer questions like these across languages. By building out our library of functions, we can enable more people to access and explore free knowledge in new ways.
Hardware devices /	Mentions of the material substrate on which Wikipedia is edited or read, factoring in hardware in decision-making.	Don't refer to image orientation such as left, right, above, or below. Image placement varies with platform and screen size, especially mobile platforms, and is meaningless to screen readers. Instead, use captions to identify images.
Imagined user	Framings of the user as designers imagine them: their skills, identity, expectations, reactions.	Lead images should be natural and appropriate representations of the topic; they should not only illustrate the topic specifically, but also be the type of image used for similar purposes in high-quality reference works, and therefore what our readers will expect to see.
Platform technical affordances /	Discussion, definition, mentions of features of Wikipedia's website, and how they fit into Wikipedia's sociotechnical system.	What are user pages for? They are not for advocacy, they are there to tell other Wikipedians about yourself, in support of collaborative editing.
Platform-independence /	References to Wikipedia data being used elsewhere, discussion of the importance of	In the case of a digital assistant or search functionality our content is integrated within a broader ontology

agnosticism	Wikipedia data to be available and in such a form that it can be used elsewhere. Non-written or visual formats using Wikipedia data.	linking multiple data sources, frequently referred to as a “knowledge graph”.
Programming / code / technology	Discussions of, documents about, policies regulating coded artefacts, activity around programming, bots, gadgets.	AFAIK, that only applies to malformed signatures that break new tools, and some rare pathological transclusion vulnerability (the latter which I don’t fully understand). It doesn’t address acceptable style, bright colours, etc.
Technical values	Values that apply mainly to, or are relevant in the context of, technological development, because they improve the functionality of a given artefact or process.	The Wikimedia Foundation created the position of steward when it recognized the need to separate user rights management from software development and systems administration.
Using technology as a debating strategy	When a feature of the platform is used to make a point / influence content production.	it is perfectly possible to engage in an edit war without breaking the three-revert rule, or even coming close to doing so.
Wikipedia’s appearance / aesthetics / design / UX	What looks good? What doesn’t? What should Wikipedia as a website look like?	I can’t remember when we last had a change to the rules about signatures, other than the WMF reform to them which has proved remarkably uncontentious for a WMF change.

Group 7 – Values

Code	Definition	Example
Accountability / transparency	Evoking the idea that editors (especially high up in the ranks) should act in public, and be available to be contacted. Also refers to process: processes are supposed to be open for anyone to see.	I am grateful to you for your openness and desire to explain the project to the widest possible range of wiki projects contributors.
Bias	Pointing to conflicts of interest, or evidence of biased reasoning in sources or in the	We need to note Alan Stern’s bias here - it appears to me that (at least to an extent) the idea of the New Horizons

	statements of other editors.	mission was sold to the US administration by pointing out that it was the only unvisited planet.
Change / lack thereof / resistance to change	Discussion of changes in the setup or policies structuring and regulating Wikipedia. Expressions of resistance to change and nostalgia.	There are definitely times where developer actions have caused sufficient harm that the Wikimedia community is up in arms; this is not an organization where "move fast and break things" works very well. Unfortunately, those relatively rare occurrences are what people remember all the time.
Community	Appeals to community as a value: the idea that the editor community has an important role to play in knowledge-production and decision-making.	Blind and absolute neutrality is for the article space. Yes, I am suggesting that we can as a community take a stand outside of the article space and should when it is not that complicated a choice.
Cooperation	Cooperation as a value: instances in which its role is discussed, in which cooperation justifies a decision or the introduction of a feature.	5. Collaborate radically; don't sign articles. Radical collaboration, in which (in principle) anyone can edit any part of anyone else's work, is one of the great innovations of the open source software movement. On Wikipedia, radical collaboration made it possible for work to move forward on all fronts at the same time, to avoid the big bottleneck that is the individual author, and to burnish articles on popular topics to a fine luster.
Copyright	Mentions of copyright, and copyright-related concepts (license, free content). As well as a legal category, copyright here is an ethical one: free content is superior to non-free content, for instance. This code applies to evidence of the role of copyright within Wikipedia's value system.	(I don't endorse this movie, if nothing else because you'll probably need to accept DRM to watch it legally)
Diversity / equity / non-Western knowledges / epistemic injustice	Moments in which issues around what knowledges are represented and how are raised. Features that support/don't support different kinds of knowledges. Inclusivity in	I assume you don't actually agree with trying to drive off editors who oppose same-sex marriage, right?

/ inclusion	general.	
Forking	Forking is the ability to make a copy of a piece of software and modify it independently. Specifically to Wikipedia, forking doesn't just involve software, but data as well: Wikipedia's content is forkable (but not personal data about editors). This code maps how forking is used in conversation, defined, applied to various parts of the project.	References that exist in a way that assumes the reader is using an encyclopedia, without reference to the specific encyclopedia (Wikipedia) or the manner of access (online), are acceptable.
FOSS	Values coming from the free and open source software (also known as FLOSS – free, libre and open source software) movement. Appeals to one value within that value system in order to support another, adjacent, value from the same value system. (eg. we are free to read, hence we should not hold copyright).	it seems to me that this project is now going in the direction of a very typical selling free software business model, namely support services, and should continue in that spirit.
Freedom / free speech	Discussion of free speech or freedom: what it is, how it manifests on Wikipedia. Uses of free speech or freedom as justification for actions or decision-making.	Generally speaking and with exceptions, individual freedom trumps the greater good at en-wiki. Custom signatures are another example, where the community consistently refuses to impose even minimal bright-line restrictions.
Individualism	Discussions of behaviour motivated by self interest or selfishness. Accusations of acting out of self-interest or selfishness. Defenses of individualism.	Of course, individual Wikipedia editors may want to set priorities for themselves. -- Toby
Law / legal justice	Discussions of the law, as is defined and enforced outside of Wikipedia. Legalistic expressions. Appeals to legal categories in conversation.	Oversighters sign confidentiality agreements and are legally required to keep the information private, while this is not true of administrators and others.

Neutrality	When neutrality is invoked as a value.	Wikipedia needs to stay neutral. Otherwise, it will not be perceived as a trusted source anymore, but as an outpost of the radical left.
Objectivity / subjectivity	Epistemic value of objectivity; accuracy. Objective vs subjective.	I echo the words of Jimbo in his talk page: I think that it's bad to allow userboxes on one side of a live political issue, while not allowing them on the other side.
Openness / ability to contribute	Evoking the ideal of Wikipedia as a project anyone can participate in, and should be able to participate in.	"Anyone can edit" is a non-negotiable principle of Wikimedia projects. Full revocation of editing permissions from any user group will not be enabled.
Practical considerations / common sense / applicability	Appeals to common sense (or related concepts, such as practicality, ease of use, and similarly intuitive justifications for action).	Pulling together wide and varied viewpoints from the communities to attempt to distill into simple actionable terms the types of behaviour that are acceptable in our collaborative environments globally is a monumental challenge.
Privacy	Mentions of privacy as a value, uses of privacy in decision-making. Justifying actions based on the defense of editors' privacy; questioning the primacy of privacy.	Another flaw in the total prohibition of "doxing" is where EditorA causes EditorB so much harm that EditorB sees fit to sue EditorA in a court of law where he can obtain financial compensation for the harm done. (Wikimedia can permanently block EditorA, but is almost powerless to prevent EditorA spawning sockpuppets and certainly cannot award EditorB damages. In order to go to court, it is necessary for EditorB to give the court EditorA's name and address which, according to Wikimedia's rules, is prohibited (See for example the fictitious example given in en:Wikipedia:Don't overlook legal threats).
Profit / money / financial	Discussions of financial matters, or the role of profit in the functioning of Wikipedia.	Non-profits should in general do what they can to avoid the problem of a single dominating revenue stream. From a governance and fiscal perspective, the

		best way to counter to that is ensuring a diverse & sustainable resource base and a healthy reserve. Diversification is something that the Wikimedia Foundation has been working to improve over much of the last decade.
Scientism / science as a referent for truth	Editors using conventions, principles, standards borrowed from science as a yardstick for knowledge-production. Scientific truth as a model/stand-in for truth.	Just like natural selection, happiness or satisfaction or quality may not be selected for, but reproductive fitness is. In Wikipedia, it may not be so much quality that is selected for, but acceptability and faddishness. The ethos within Wikipedia is that the ecology of Wikipedia is self-healing and does indeed breed quality. What this concept suggests is that it is important to design the software, and set policy, such that the ecology of Wikipedia does indeed breed for quality. Currently, it seems generally accepted that it does.
Universality / ubiquity	Mention of Wikipedia's ambition to being universal, disseminating knowledge globally. Also universal in content – covering the whole of human knowledge. Both reach and content.	Our donors want us to serve more people in more parts of the world and understand that that requires revenue growth
Wikipedia's mission / purpose	Definitions of Wikipedia's aim, or uses of one's idea of what Wikipedia's aims are in order to support a statement or justify behaviour.	We are not here to provide nice neat edit summaries or to make our signatures look good, but to build an encyclopedia.