

A Critical Take on the Policy Recommendations of the EU High-Level Expert Group on Artificial Intelligence

Michael Veale

Faculty of Laws, University College London and the Alan Turing Institute

October 2019

Forthcoming in the European Journal of Risk Regulation

The European Commission recently published the policy recommendations of its 'High-Level Expert Group on Artificial Intelligence': a heavily anticipated document, particularly in the context of the stated ambition of the new Commission President to regulate in that area. This essay argues that these recommendations have significant deficits in a range of areas. It analyses a selection of the Group's proposals in context of the governance of artificial intelligence more broadly, focussing on issues of framing, representation and expertise, and on the lack of acknowledgement of key issues of power and infrastructure underpinning modern information economies and practices of optimisation.

1	Background	1
2	Framing Issues	3
3	Low Level Experts Needed	4
4	Lack of Infrastructure	5
5	Power Outage	7
6	Concluding Remarks	8

1 Background

It is currently trendy to establish bodies to ponder upon the topic of artificial intelligence (AI) for policymakers. In the United Kingdom, this has been taken to a bizarre extreme. Attached to the public sector alone we can identify at least: the Centre for Data Ethics and Innovation;¹ the AI Council;² the Office for AI;³ the House of Lords Select Committee on AI;⁴ the House of Commons Inquiry on Algorithms in Decision-Making;⁵ the Alan Turing Institute; the National Data Guardian;⁶ the Information Commissioner's Office; a proposed new digital regulator;⁷ the various directorates of the Department of Digital, Culture, Media and Sport who took over from the Government Digital Service on draft-

¹Department of Digital, Culture, Media and Sport, *Centre for Data Ethics and Innovation: Government Response to Consultation* (UK Government 2018).

²Leading Experts Appointed to AI Council to Supercharge the UK's Artificial Intelligence Sector' (GOV.UK, 16 May 2019) <<https://www.gov.uk/government/news/leading-experts-appointed-to-ai-council-to-supercharge-the-uks-artificial-intelligence-sector>> accessed 15 August 2019.

³Office for Artificial Intelligence, 'About Us' (GOV.UK, 2019) <<https://www.gov.uk/government/organisations/office-for-artificial-intelligence/about>> accessed 15 August 2019.

⁴House of Lords, 'AI in the UK: Ready, Willing and Able?' (HL Paper 100, 16 April 2018).

⁵House of Commons, 'Algorithms in Decision-Making' (HC 351, 23 May 2018).

⁶Health and Social Care (National Data Guardian) Act 2018.

⁷HM Government, 'Online Harms White Paper' (CP 57, April 2019).

ing the Government Data (Science) Ethics Framework;⁸ the Digital Catapult;⁹ the Office for Tackling Injustices;¹⁰ the Regulatory Horizons Council;¹¹ proposed new algorithmic accountability powers against platforms for Ofcom;¹² NHSX, a new health National AI Lab;¹³ AI monitoring by the Health and Safety Executive's (HSE) Foresight Centre;¹⁴ AI analysis from the Government Office for Science;¹⁵ the Office for National Statistics' Data Science Campus;¹⁶ and the Department of Health and Social Care's code of conduct for data-driven health and care technology.¹⁷ Keeping up with these organisations, initiatives and efforts is a full-time job in and of itself.

The European Commission has not set up quite so many initiatives as of yet, but among those they have, the most well-known is arguably the 'High-Level Expert Group on Artificial Intelligence'.¹⁸ This group was formed through individuals nominating themselves as experts in this field, and established in 2018. In the year following its establishment, it has produced two main deliverables. The first was a set of ethics guidelines.¹⁹ These ethics guidelines did not land, even in their revised post-consultation form, without controversy and accusations of 'ethics washing' and industry capture of the process.²⁰ The second deliverable was a set of 'policy and investment' recommendations.²¹ This article critiques that document, and contextualises that critique in a broader narrative about the trajectories of artificial intelligence, policy and society.

This document does come at a time of agenda-setting in Europe. In May 2019 the European Commission announced its intention to create 'robust regulatory framework [to] proactively address the ethical and legal questions surrounding artificial intelligence.'²² This was re-iterated in Commission President-elect Ursula von der Leyen's 'Agenda for Europe', where she stated her intention to 'put forward legislation for a coordinated European approach on the human and ethical implications of Artificial Intelligence'.²³ It is therefore imperative that its recommendations should be treated to maximum possible scrutiny.

This essay focusses on several issues with the HLEG's policy recommendations in the area of artificial intelligence. Firstly, the document is set up as to ignore and encourage the skipping of problem framing in areas where AI is suggested as being helpful. The capacities needed to assess the usefulness of computing in any given context, rather than just to increase its availability or intensity, go unconsidered. Secondly, the type of

⁸Department of Digital, Culture, Media and Sport, 'Data Ethics Framework' (GOV.UK, 30 August 2018) <<https://www.gov.uk/government/publications/data-ethics-framework/data-ethics-framework>> accessed 15 August 2019.

⁹Digital Catapult, 'Ethics Framework — Responsible AI' (MI Garage, no date) <<https://www.migarage.ai/ethics-framework/>> accessed 15 August 2019.

¹⁰HM Government, 'PM Announces New Independent Organisation to Tackle Deep-Rooted Injustices in Society' (GOV.UK, 12 July 2019) <<https://www.gov.uk/government/news/pm-announces-new-independent-organisation-to-tackle-deep-rooted-injustices-in-society>> accessed 15 August 2019.

¹¹Department for Business Energy & Industrial Strategy, 'Regulation for the Fourth Industrial Revolution' (CP 111, June 2019).

¹²HM Government (n 7).

¹³Tara Donnelly and Sam Roberts, 'Introducing NHSX's New National Artificial Intelligence Laboratory' (*Technology in the NHS (Department of Health & Social Care Blog)*, 8 August 2019) <<https://healthtech.blog.gov.uk/2019/08/08/introducing-nhsx-new-national-artificial-intelligence-laboratory/>> accessed 15 August 2019.

¹⁴HL Deb 5 June 2019, WA HL8200.

¹⁵Government Office for Science, 'Artificial Intelligence: Opportunities and Implications for the Future of Decision Making' (GO Science Report, 2016).

¹⁶Office for National Statistics, 'About Us: Data Science Campus' (no date) <<https://datasciencecampus.ons.gov.uk/about-us/>> accessed 15 August 2019.

¹⁷Department of Health & Social Care, 'Code of Conduct for Data-Driven Health and Care Technology' (GOV.UK, no date) <<https://www.gov.uk/government/publications/code-of-conduct-for-data-driven-health-and-care-technology/initial-code-of-conduct-for-data-driven-health-and-care-technology>> accessed 15 August 2019.

¹⁸This arose from an earlier communication, see European Commission, 'Artificial Intelligence for Europe' (Communication from the Commission (SWD(2018) 137 final), April 2018).

¹⁹High-Level Expert Group on Artificial Intelligence, 'Ethics Guidelines for Trustworthy AI' (April 2019).

²⁰Thomas Metzinger, 'Ethics Washing Made in Europe', *Der Tagesspiegel* (18 April 2019) <<https://www.tagesspiegel.de/politik/eu-guidelines-ethics-washing-made-in-europe/24195496.html>> accessed 15 August 2019; 'Dialogue Seminar on Artificial Intelligence: Ethical Concerns' (19 March 2019) <<http://www.europarl.europa.eu/streaming/?event=20190319-1500-SPECIAL-SEMINAR1&start=2019-03-19T15:44:53Z&end=2019-03-19T15:56:00Z&language=en>> accessed 15 August 2019.

²¹High-Level Expert Group on Artificial Intelligence, 'Policy and Investment Recommendations for Trustworthy AI' (26 June 2019).

²²European Commission, 'Europe in May 2019: Preparing for a More United, Stronger and More Democratic Union in an Increasingly Uncertain World' (Contribution to the informal EU27 leaders' meeting in Sibiu (Romania) on 9 May 2019, 9 May 2019) 33.

²³Ursula von der Leyen, 'A Union that Strives for More: My Agenda for Europe' (Political Guidelines for the Next European Commission 2019-2024, 2019) 13.

expertise and input recommended to cope with this issue — largely an increased, inward-looking focus on ethicists of artificial intelligence or stakeholders in an ‘AI’ bubble — do not seem appropriate responses to factors and environments that foster unethical technology in practice today, and may even further marginalise the voices of affected communities. The structure and the composition of the HLEG only contributes to this problem. Lastly, the HLEG either ignore or under-emphasise two critical issues: the infrastructure underpinning optimisation using ‘AI’, and questions of power and control.

2 Framing Issues

What is meant when organisations apply ‘AI’ to a problem is often indistinguishable from the application of computing, statistics, or even evidence. The usage of the phrase has become so laughably ambiguous and general, it is almost like saying that to solve an urban infrastructure problem, one must ‘apply power tools’. Yet, unlike that analogy, ‘AI’ is ‘intelligent’. This masks an important way in which no tool can be intelligent: tools cannot define the problem they are applied to. When holding a hammer, everything looks like a nail.

Because of this, a focus on AI incentivises fudging or skipping arguably the most important stage in any policy process — problem structuring and definition.²⁴ Take the HLEG’s recommendation to ‘[f]oster the development of AI solutions that address sustainability challenges.’²⁵ There is no blueprint for engineering societal sustainability. It is widely acknowledged as a ‘wicked’, poorly structured problem. There is limited political agreement on the *ends* of sustainability — who will sacrifice what for carbon reductions? — and further limited agreement on the *means* by which these disputed ends should be sought.²⁶ Deploying artificial intelligence as a ‘solution’ requires quite certain agreement on ends, which are incidentally limited to those compatible with rendering the world legible by machines.²⁷ The difficulty in reaching agreement on ends does not, however, result from general dithering and indecision which a keen (and likely white cis male) engineer can efficiently resolve, coming up with ends suitably compatible to optimisation through ‘AI’. Agreement can only be sustainably reached by genuine processes of negotiation and participation.²⁸ Bulldozing through or subverting that stage is not a recipe for sustainability, but for tension, conflict, and unintended consequences. The idea that exists as a subtext that AI is ‘intelligent’ enough to abolish trade-offs, depoliticise extremely contested issues, and give all parties what they seek is at best unevidenced and naïve, and at worst a lie.

The HLEG recommendations, among many policy reports and views in this area, suffer from a deficit in recognising the importance of problem structuring and framing more broadly. We can take some lessons on how best to apply the ‘magic’ of artificial intelligence from the application of actual magic by late fantasy author Terry Pratchett. Pratchett’s popular, comical *Discworld* series contains two main users of magic: wizards and witches. Bureaucratic, hierarchical wizards blindly frame all problems as soluble by magic, often bringing problematic consequences. Witches, who live independently in superstitious rural villages, see the most important feature of magic as knowing *when not to use it*. They are rarely seen casting spells, as they start with the problem in front of them, often passing off psychology, practical knowledge or social dynamics as ‘magic’ to achieve their ends. The HLEG, in many recommendations, emphasises training and retaining AI specialists for the application of AI to societal problems. It misses the importance of having the capacity to *know when not to use artificial intelligence. To what extent is ‘intelligent’*

²⁴See generally Eugene Bardach, *A Practical Guide for Policy Analysis: The Eightfold Path to More Effective Problem Solving* (Sage 2012) (on the policy process).

²⁵High-Level Expert Group on Artificial Intelligence (n 21) 12.

²⁶See generally Robert Hoppe, *The Governance of Problems: Puzzling, Powering and Participation* (The Policy Press 2010) (on the nature of unstructured problems and their management by the policy process).

²⁷On the limitations of machine-legible approaches to complex problems, see generally James C Scott, *Seeing like a State: How Certain Schemes to Improve the Human Condition Have Failed* (Yale University Press 1998); Sally Engle Merry, *The Seductions of Quantification: Measuring Human Rights, Gender Violence, and Sex Trafficking* (University of Chicago Press 2016).

²⁸Hoppe (n 26).

surveillance a solution for anti-social youth behaviour in an area without enough alternative evening activities or public spaces? What use is predictive railway maintenance if the infrastructure was poorly specified and invested in further upstream? Will monitoring and prediction of employee workload in a search for productivity compensate for the mental stress of constant observation such a system might engender? Ben Green addresses these issues in his recent book *The Smart Enough City*, highlighting the need to see technology as just one tool in a toolkit: just one of many means that might potentially be used to head towards a complex and societally negotiated end.²⁹ The focus should be on taking off the 'tech goggles' to identify problems, challenges and needs, and to not be afraid to discover that other policy options are superior to a technology investment.

The HLEG recommendations are, at times, aiming at deficits that do not exist. Unscrupulous vendors do not need help to better locate businesses and public bodies apparently 'in need' of an AI 'solution'. Their PR efforts are doing a stellar job in that department as it stands. On the contrary, intervention might be needed in helping buyers (public and private) lacking in analytic capacity figure out whether there is truly a promising role for technology in their problems, and what it might look like. Those inside organisations — procurement officials, advisors, policy officers, and the like — must be able to cut through the advertising and lay out options in a fair and even way, placing technologies such as artificial intelligence alongside, not above, the many other tools at their disposal.

3 Low Level Experts Needed

A similar issue can be taken with the surprising suggestion the report supports that '720 professorial chairs in applied ethics' are required across a network of 'trustworthy AI' research.³⁰ Is a lack of institutionalised ethicists the source of the problems being faced today? Or, more cynically, might this suggestion be a result of the interests of those sitting on the HLEG itself?

As the *Tech Won't Build it* and *Google Walkout* movements are illustrating, organisational factors permitting the unethical use of technology seems to be located in the ways the business models of tech giants override the desires and concerns of individual expert workers. It would be quaint if the main or only barrier to reflection and change for a lumbering, bureaucratic multinational was the lack of a well-meaning ethicist to point out a quandary or unintended consequence. Quaint, but unrealistic. The business models that scale to create today's multinational money-printing (or, at the very least, capital-attracting) machines generate externalities which are just the cost of doing business. The harms, risks and inconveniences they shift to society are in general inextricable parts of the business model, with some more amenable to cheap mitigation (often begrudgingly in the face of potential or actual, enforced legislation) than others.

If movements like *Tech Won't Build It* and *Google Walkout* have taught anything, it's that there is some, perhaps limited, hope in the moral of individual workers. Yet are ethicists the solution here? They seem extraordinarily poorly placed to train workers who are skilled and empowered to organise, object and influence business models from within — indeed, a training in *assessing* ethical issues says little about how to action the results of those assessments. What is needed is a challenge to corporate governance. Yet any such reforms which might shake-up corporate decision-making or board representation are conspicuously *completely absent* from HLEG recommendations.

An issue with *both* ethicists and employees, largely engineers, taking the lead surrounds the flaws in an engineering-centred approach to ethical issues. Seeta Peña Gangadharan and Niklas Jędrzej have recently critiqued the techno-centric nature of studies of fairness, accountability and transparency, drawing attention to how civil society voices analyse technology as part of systemic practices of injustice, rather than miss the role of broader

²⁹Ben Green, *The Smart Enough City* (MIT Press 2019).

³⁰High-Level Expert Group on Artificial Intelligence (n 21) 24.

systems.³¹ Seeing a greater role for ethicists of technology or conscientious engineers as the correct response to injustice exacerbated by technology but not, at its root, caused by it risks further marginalising those with the clearest view of on-the-ground issues and the closest connection and legitimacy to affected communities. The HLEG is notable by its exclusion of such voices, seeing expertise in artificial intelligence as the domain of technical researchers, generalist ethics and governance scholars, industry lobby groups. This is heavily problematic, particularly in light of the extensive research on the factors for success of expert assessments in value-laden, technically challenging areas, which illustrate the importance of legitimacy in the success of these assessments.³² Affected communities ‘judge legitimacy [of expert assessments] based on who participated and who did not, the processes for making choices, and how information was produced, vetted, and disseminated’, with ‘[p]otential users often reject[ing] assessments that lack legitimacy as “not invented here”’.³³ As it stands, only two of the fifty-one members of the HLEG clearly represent civil society of *any* flavour, let alone in any of the specific domains the report discusses.³⁴

4 Lack of Infrastructure

The HLEG recommendations are additionally deeply timid in respect to even discussion of the need to exercise of legal control over data collection infrastructures.

The capabilities of the largest firms to build effective machine learning models is often more of a result of the data collection and experimentation *infrastructure* they have established than any single static dataset. Signs of these infrastructures are everywhere, from the data-intensive exhaust of cloud services powering voice assistants or location services to the covert third-party trackers littering the Web³⁵ and popular apps.³⁶ They connect to the new way software is designed, through ‘agile’ methods in ‘perpetual beta’, where the invisible experimentation, data collection and updating infrastructures mean that for many pieces of software the version you use today may not be the same as that of the person sitting beside you, and is unlikely to be the same version you use next week.³⁷ These firms can build machine learning models of all flavours with this infrastructure, affording them informational power in a huge range of situations rather than just a handful.

Equally important is the role of infrastructure in deploying artificial intelligence. Models matter most when their results are used, and so those with the power to propagate the results of these models at scale are just as critical as points of policy intervention as those where models are created and trained. The report does not adequately address the notion that in some cases, artificial intelligence may change the delicate, societally negotiated power balances relating to existing technical infrastructures, such as financial institutions and online intermediaries. The current societal acceptance of CCTV cameras among some individuals could be argued to have been partly based on the notion that the footage was sufficiently labour intensive to examine and monitor such that it would not be considered constant surveillance. The application of machine learning techniques, such as machine vision, and the use of automatic classifiers such that CCTV is ‘always being watched’ changes the power of those controlling previously ‘dumb’ infrastructure.

³¹Seeta Peña Gangadharan and Jędrzej Niklas, ‘Decentering Technology in Discourse on Discrimination’ (2019) 22 *Information, Communication & Society* 882.

³²See generally the chapters (and works cited therein) in the collection Ronald B Mitchell and others, *Global Environmental Assessments* (MIT Press 2006).

³³Ronald B Mitchell and others, ‘Information and Influence’ in Ronald B Mitchell and others (eds), *Global Environmental Assessments* (MIT Press 2006) 321.

³⁴These are representatives from Access Now, a digital rights organisation, and BEUC, an umbrella body for consumer rights groups.

³⁵Zhonghao Yu and others, ‘Tracking the Trackers’ in *Proceedings of the 25th International Conference on World Wide Web (WWW ’16)*, Republic and Canton of Geneva, Switzerland, International World Wide Web Conferences Steering Committee 2016).

³⁶Reuben Binns and others, ‘Third Party Tracking in the Mobile Ecosystem’ in *Proceedings of the 10th ACM Conference on Web Science (WebSci ’18)*, New York, NY, USA, ACM 2018).

³⁷Seda Gürses and Joris van Hoboken, ‘Privacy after the Agile Turn’ in Evan Selinger and others (eds), *The Cambridge Handbook of Consumer Privacy* (1st edn, Cambridge University Press 2018).

In these cases, the use of AI might require existing policies to be reconsidered with a view to retaining the same balance of societal interests and fundamental rights, which might mean limiting infrastructural expansion that previously seemed proportionate.

In general, tackling the policy challenges of AI cannot be divorced from tackling the policy challenges of tracking and sensing, both in digital and physical infrastructures. To this extent, it is disappointing that the HLEG do not discuss the lack of regulatory capacity to analyse, proactively monitor and police these areas.

The Information Commissioner's Office (ICO), the UK's data protection supervisory authority, is the largest in Europe. It receives no money from the treasury, and its funding comes almost entirely from fees levied on data controllers, with a budget of approximately £40m.³⁸ Given that all data controllers fall in scope as being subject to regulation, as well as the 66m inhabitants of the United Kingdom and any non UK residents anywhere in the world who are data subjects in relation to a UK data controller, this is a huge undertaking for that sum of money. Other regulators of digital services, such as Ofcom, the broadcast regulator, have a budget nearly five times greater, and often spend above it.³⁹ The ICO is furthermore in strong shape in terms of number of employees compared to other European data protection authorities. A recent report from the European Data Protection Board paints a dismal picture of supervisory authorities' financial and human resourcing needs, with several regulators in the EU stating their need for 50% more or even double the resources they currently can muster.⁴⁰

Without sufficient regulatory resource, it is not going to be possible to understand, monitor and assess the infrastructures that artificial intelligence transforms and it used through in a way that can truly ensure they comply with the law. As it stands, infrastructure such as real-time bidding, an advertising system drawing heavily on machine learning-powered profiling that has been powering a significant proportion of adverts on the web since 2010, has only just been analysed and found to be significantly illegal by the Information Commissioner.⁴¹ The illegality of this pervasive system is not new, but a ten year wait for analysis is simply too long to wait to uncover significant areas of regulatory non-compliance.

The HLEG do not recommend increased regulatory and enforcement capacity. The main public capacity recommendation the report *does* provide surrounds a perceived need for data sharing infrastructures and open public data (which appear strangely ignorant of the 2019 Open Data Directive, its 'open by design and by default' public sector provisions⁴² and particular requirements for 'dynamic data' and 'high-value' data being made available through APIs and bulk downloads). The HLEG should have instead taken the opportunity to recommend an increase in the funding and capacity for all *regulators* in light of technological challenges. The proper funding of regulators is a primary inhibitor of the enforcement of existing laws today, and the even and proper application of laws across the Single Market, without *de facto* regulatory arbitrage, relies on at least comparable capacity between Member States. Again, a cynic might note the presence of several large firms and notorious technology lobby groups such as DIGITALEUROPE in the HLEG as a reason why increased enforcement of commonly-flouted existing data law did not make it into the report.

³⁸Information Commissioner's Office, 'Information Commissioner's Annual Report and Financial Statements 2018-2019' (Report presented to Parliament (HC 2299), 8 July 2019) 98.

³⁹Ofcom, 'The Office of Communications Annual Report & Accounts for the Period 1 April 2018 to 31 March 2019' (Report presented to Parliament (HC 2321), 2019) 103.

⁴⁰European Data Protection Board, 'First Overview on the Implementation of the GDPR and the Roles and Means of the National Supervisory Authorities' (Report presented to the European Parliament's Civil Liberties, Justice and Home Affairs Committee (LIBE), 26 February 2019).

⁴¹Information Commissioner's Office, 'Update Report into Adtech and Real Time Bidding' (20 June 2019).

⁴²Directive (EU) 2019/1024 of the European Parliament and of the Council of 20 June 2019 on open data and the re-use of public sector information OJ L 172/56 (2019) art 5(2).

5 Power Outage

Power rarely makes a bold showing in the recommendations of the HLEG. Indeed, the few places where HLEG have mentioned redressing relative power and control into the document are easily identified by both being placed at the end of their respective sections, and by beginning with a non-committal recommendation to ‘consider’ instead of ‘develop’, ‘set up’ or ‘adopt’. The few that exist are almost all weak or ambitious in their content, too. For example, the public sector should ‘consider’ mandating their contractors to return ‘data’ (not models) produced in the context of a contract to a public body paying them — unless of course it was not of sufficient ‘general public interest’ and/or it might ‘infringe the private company’s IP’.⁴³

Recommendation 18.8 is one of the only daring recommendations in the entire document in this regard, suggesting that legally mandated interoperability requirements should be imposed on a sector-by-sector basis on both dominant and non-dominant firms to enable competitive market infrastructures.⁴⁴ Such recommendations are far from new,⁴⁵ but constantly shot down by firms who much prefer the framing of ‘data portability’ (where a non-dynamic snapshot can be moved from provider-to-provider by a particularly determined consumer). The dominant players in AI are currently collaborating to ensure they can pre-empt legislative desires and point to functioning snapshot-style data portability, rather than risk seeing interoperability on the regulatory agenda.⁴⁶

The only other recommendation concerning competition and market power is to ‘consider’ revisions to competition law to consider ‘the volume of data or incumbency data advantages – the building block of many AI systems – in the assessment of market power’.⁴⁷ This recommendation sounds worthy, and certainly would be impactful, but is already one step behind the current trajectory of the largest technology firms in the world.

To understand why, it is important to note that firms do not intrinsically care about data, but their *ability to optimise*.⁴⁸ Governments and surveillance agencies are the main candidates for caring about the content of datasets beyond their optimisation potential, but most private firms operating with broadly lawful intentions do not have an interest in qualitatively snoop on individuals for their own business purposes. Typically, it has gone without saying that you cannot optimise without data. This left large platforms heavily vulnerable to co-optation by state surveillance, as the leaked Snowden documents illustrate.⁴⁹

Today, a growing number of firms have taken steps to publicly tie their hands, practically and mathematically limiting their ability to view individual records of data *whilst still* controlling infrastructure, often cryptographic or hardware-based in nature, which allows them to run the optimisations or deliver the services that bring in revenue. The most high-profile case of this is the well-publicised struggle between Apple and the FBI in relation to accessing the contents of an encrypted iPhone used in a terrorist shooting at San Bernadino, CA in late 2015.⁵⁰

Many technologies are being fast developed that permit optimisation without a firm being able to see individual records of data ‘in the clear’— in decrypted form. These include technologies which allow aggregate statistics or machine learning models to be computed without the pooling of data in a single place, or beyond devices (such as homomorphic encryption or secure multi-party computation) and techniques designed to establish datasets which do not reveal information about individual users (such as dif-

⁴³High-Level Expert Group on Artificial Intelligence (n 21) 19–20.

⁴⁴*ibid* 29–30.

⁴⁵See e.g. Ian Brown and Christopher T Marsden, ‘Holistic Regulation of the Interoperable Internet’ in *Regulating code: Good Governance and Better Regulation in the Information Age* (The MIT Press 2013).

⁴⁶See the effort by Facebook, Google, Microsoft, Apple and Twitter: Data Transfer Project, ‘Data Transfer Project Overview and Fundamentals’ (White Paper, 20 July 2018).

⁴⁷High-Level Expert Group on Artificial Intelligence (n 21) 39.

⁴⁸Seda Gürses and others, ‘Stirring the POTs: Protective Optimization Technologies’ in Emre Bayamlioğlu and others (eds), *Being Profiled: Cogitas Ergo Sum* (Amsterdam University Press 2018).

⁴⁹Caspar Bowden, ‘The US Surveillance Programmes and Their Impact on EU Citizens’ Fundamental Rights’ (Report for the European Parliament’s Committee on Civil Liberties, Justice and Home Affairs, 2013).

⁵⁰Karl Stephan, ‘Apple Versus the Feds: How a Smartphone Stymied the FBI’ (2017) 6 *IEEE Consumer Electronics Mag* 103.

ferential privacy).⁵¹ These are good from the standpoint of privacy-as-confidentiality, for individuals whose threat models consist of qualitative information about themselves being leaked. They are not as good for ensuring individual or society control concerning design or optimisation decisions, particularly those designed to profile them, manipulate them, or to place different weightings or disregard aspects of society or the environment.⁵² Those concerned about data-driven ‘nudging’ as a mode of regulation⁵³ will likely find little solace in the fact that confidentiality of data is mathematically assured despite firms be able to deploy ‘artificial intelligence’ to influence individuals in effectively the same way as previously.

6 Concluding Remarks

The HLEG missed a range of opportunities in this report. The document is a marathon of recommendation, with no clear prioritisation and the more interesting of the recommendations buried deeply in a range of futile, confused or non-actionable suggestions. It foregrounds AI as a technological solution to completely inappropriate issues without considering the capacities needed to understand the problems which AI could potentially be applied to. It incorrectly portrays the reasons why unethical technologies are found used in the wild today, making them appear like *bone fide* oversights that ethicists might be able to highlight rather than intrinsic parts of business models which disregard their effects on societies and environments. By their composition and their recommendation, they further side-line and marginalise community and domain voices, and seek to reify an elite club of AI and society experts to the detriment of those with connection to harms and issues that technologies exacerbate. And finally, they largely ignore infrastructure and power, which should be one of, if not the most, central concern around the regulation and governance of data, optimisation and ‘artificial intelligence’ in Europe going forwards.

There is limited time in which to shape tomorrow’s information infrastructures while they are still, even slightly, in flux. The European Commission needs challenging thinkers that are bold and perceptive in their understanding of the future trajectories of those firms intermediating so much of today’s information flows. Unfortunately, as a collective and in my opinion, the High Level Group has not demonstrated that they can provide that critical analysis.

Acknowledgements

Thank you to Seda Gürses for her responsive term ‘low-level experts’, and for thinking through many of these issues in relation to our joint consultation response to the earlier HLEG-AI guidance. Thank you to the anonymous reviewer whose input improved this work. Michael Veale was supported by the Alan Turing Institute under EPSRC grant EP/N510129/1.

⁵¹See generally Josep Domingo-Ferrer and others, *Privacy and Data Protection by Design—from Policy to Engineering*. (ENISA 2014); The Royal Society, *Protecting Privacy in Practice: The Current Use, Development and Limits of Privacy Enhancing Technologies in Data Analysis* (The Royal Society 2019).

⁵²Gürses and others (n 48).

⁵³See, e.g., Karen Yeung, ‘“Hypernudge”: Big Data as a Mode of Regulation by Design’ (2017) 20 *Information, Communication & Society* 118.