

# Design Digital Multisensory Textile Experiences

Shu Zhong  
University College London  
London, United Kingdom  
shu.zhong.21@ucl.ac.uk

## Abstract

The rise of Machine Learning (ML) is gradually digitalizing and reshaping the fashion industry, which is under pressure to achieve Net Zero. However, the integration of ML/AI for sustainable and circular practices remains limited due to a lack of domain-specific knowledge and data. My doctoral research aims to bridge this gap by designing digital multisensory textile experiences that enhance the understanding of the textile domain for both AI systems and humans. To this end, I develop TextileNet, the first fashion dataset using textile taxonomies for textile materials identification and classification via computer vision, and TextileBot, a domain-specific conversational agent. TextileBot integrates textile taxonomies with large language models (LLMs) to engage consumers in sustainable practices. Additionally, my research explores how multisensory experiences can improve user understanding and how AI perceives textiles. The overarching goal is to embed human expertise into machines, design immersive multisensory experiences, and facilitate natural human-AI interactions that promote sustainable practices.

## CCS Concepts

• **Human-centered computing**; • **Computing methodologies** → **Artificial intelligence**; **Machine learning**; *Natural language processing*;

## Keywords

Human-AI interaction, AI for Social Good, Sustainability, Machine Learning, Multimodal Large Language Models, Agents

## ACM Reference Format:

Shu Zhong. 2024. Design Digital Multisensory Textile Experiences. In *INTERNATIONAL CONFERENCE ON MULTIMODAL INTERACTION (ICMI '24)*, November 04–08, 2024, San Jose, Costa Rica. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3678957.3688621>

## 1 Introduction

Clothing and textiles are ubiquitous in our daily lives. However, the textiles industry contributes largely to global carbon emissions. Take the UK as an example: its textile industry's CO<sub>2</sub> emissions are nearly equivalent to those of private cars [28]. Millions of tons of garments end up in landfills every year [27].

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ICMI '24, November 04–08, 2024, San Jose, Costa Rica

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0462-8/24/11

<https://doi.org/10.1145/3678957.3688621>

In response, the fashion industry—including brands and retailers—has increasingly embraced the need for creating more *sustainable fashion* [18]. A key concept in this shift is textile circularity, a sub-area of the broader circular economy that aims to transform the traditional "take-make-dispose" model. Textile circularity focuses on extending the lifecycle of textile products through reuse, repair, recycling, and upcycling, thereby reducing resource consumption and waste [11, 15]. However, implementing textile circularity can be complex and abstract, often requiring extensive domain-specific knowledge in areas such as materials science and supply chain management. Additionally, it necessitates efforts to help the general public understand and engage in these sustainable practices.

Given these complexities, there is growing interest in leveraging Artificial Intelligence (AI), including Machine Learning (ML) and Large Language Models (LLMs), to tackle specific challenges. AI is already being used in various sectors, including fashion, where it helps consumers find products that meet their needs [6]. For instance, Liu *et al.* [23] demonstrated that neural networks can effectively classify various garment attributes, such as categories, textures, and shapes. However, the application of AI in sustainable practices, such as textile recycling, is limited by the scarcity of domain-specific data for specialized tasks. This data should ideally be developed in collaboration with domain experts, not solely by computer scientists. My multidisciplinary background and involvement in the cross-disciplinary UKRI Interdisciplinary Circular Economy Centre for Textiles Circularity project<sup>1</sup> greatly inspire my research motivation.

Additionally, enhancing public engagement requires transitioning from passive consumers to active users [25]. Designed multisensory experiences can provide users with memorable interactions [34], making abstract concepts more tangible and easier to understand [12]. My research aims to craft digital multisensory textile experiences that improve both AI understanding and public involvement in sustainable practices. By integrating material science, computer science, and human-computer interaction (HCI), human-AI Interaction serves as a vital tool for conveying these concepts.

The aim of this doctoral research is to explore how AI can create digital multisensory textile experiences—incorporating vision, touch, and olfactory elements—to bridge the gap between material science and computer science. This work focuses on advancing textile circularity by integrating AI with multisensory technologies, thereby connecting consumers to sustainable practices. Human-Centred Design (HCD) [29] serves as the foundation of this work, employing an iterative process that includes empathizing, defining, ideating, prototyping, testing, and iterating.

The motivation behind this research is to investigate how AI can enhance the understanding and communication of complex,

---

<sup>1</sup>The Textiles Circularity Centre (TCC) <https://textilescircularity.rca.ac.uk/>

domain-specific concepts like textile circularity. By developing advanced AI tools and interactive experiences, the study seeks to address the following key questions:

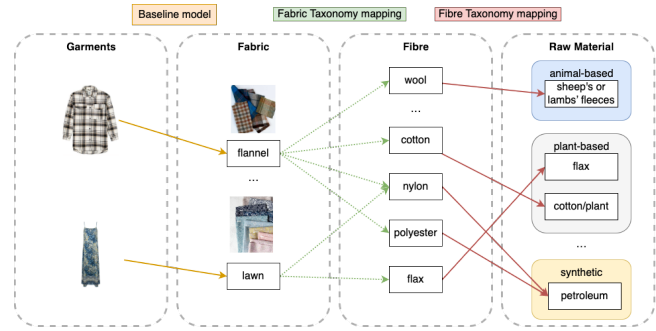
- How can AI accurately process and interpret textiles through sensory inputs such as vision, touch, and smell?
- How can we design and implement tools that effectively communicate domain-specific knowledge, such as textile circularity, to consumers?
- How can AI-driven tools assist in educating consumers about complex, domain-specific topics?
- How do consumers perceive and interact with these AI-driven tools?

## 2 Contributions to date

In the initial stage of my PhD, as an interdisciplinary research spinning across computer science and material science was crucial to gain a comprehensive understanding of the elements in textiles. The primary objective was to identify the gaps necessary to create a multisensory textile experience. By bridging the knowledge gap and creating a shared language between material science and computer science, we developed TextileNet [38]. TextileNet is the first fashion textile dataset with the most diverse labels, enabling AI to "see" and distinguish different textiles. This dataset is fundamental for defining and classifying textiles, which is essential for designing textile experiences (see Section 3).

Building on this foundation, we aimed to build a tool that can communicate about textiles circularity – let AI talk about textiles. To this end, we introduced TextileBot [37], an LLM-mediated voice conversational agent integrated with the knowledge of the TextileNet taxonomy [38]. We proposed to use taxonomy-based prompts to transform LLMs from task-agnosticism to domain-focused. To assess the effectiveness of our approach and understand the Human-AI interaction, we conducted a user study with TextileBot. We found that TextileBot can create experiences around textiles by engaging users in a conversation around textiles through our approach. We also provide some insights for designing LLM-mediated voice-based conversational agents. This work is detailed in Section 4.

Additionally, we strive to understand consumers' textile experiences from a perception perspective. Amidst a world filled with various sensory stimuli such as sounds, smells, and sights, we often take these sensory experiences for granted. We believe in the importance of cultivating mindfulness and fully embracing these sensory experiences. We investigated how AI, especially LLMs, understand textiles touch experiences [36]. Mindfulness meditation has been shown to enhance present-moment awareness and non-judgmental observation [16, 22]. To investigate whether mindfulness meditation can improve our sensory experiences, we conducted a study evaluating the effects of mindfulness meditation on mid-air haptic experiences through a recognition task on the palm. The findings indicate that mindfulness meditation can enhance the recognition performance of mid-air touch experiences [33]. We then explore using mid-air haptic to render subtle stimuli on the palm during meditation to understand whether one's sense of touch can facilitate meditation [7]. Our results suggest mid-air haptics initially distracted participants but eventually facilitated mindfulness through three processes.



**Figure 1:** We present two garment examples in TextileNet-fabric to predict fabric labels (*flannel* and *lawn*), and show how they can use our taxonomies to create relations to different fibres and raw materials. For instance, for *lawn*, we can connect it to *nylon* and *flax*, this then further connects us to the raw material space with *flax* and *petroleum*.

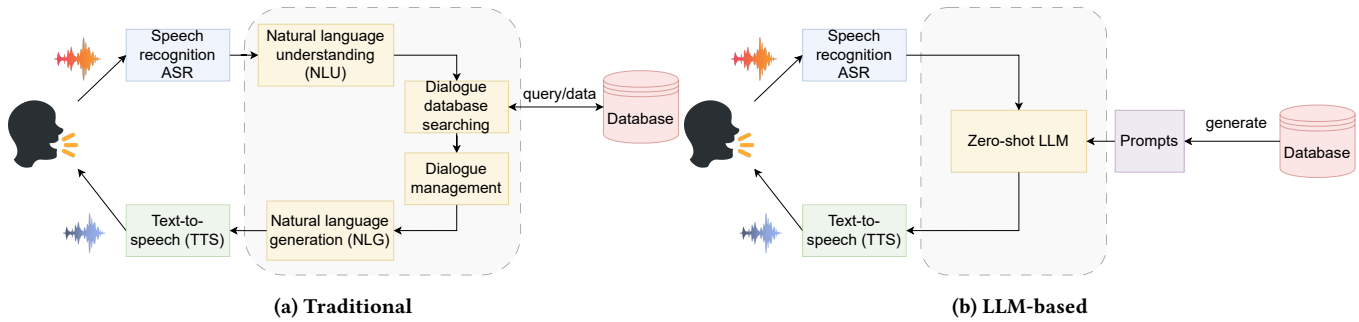
## 3 AI can "see" the differences between textiles – TextileNet

Yet, nowadays, textiles are mostly sorted manually [30], despite recent research raised using near-infrared spectroscopy (NIR spectroscopy) to recognize textiles for automated garments sorting line [9]. A low-cost, high-efficiency technique for automatically identifying textile materials in garments is missing. Such techniques could enable the digitized fashion sector to retrieve the materials in garments, promoting automated recycling, thereby significantly reducing textile waste and carbon emissions [10].

Given the importance of textile material identification in clothing, it's crucial to note that this process can be complex due to the distinction between fibre and fabric materials. Fibre is the raw material used to make fabric, but most existing fashion datasets in machine learning mix them in the same category. Currently, no dataset contains organized textile material labels reviewed by materials scientists [14, 23, 39].

In light of this, we built the first fashion textile dataset, TextileNet, based on textile material taxonomies - a fibre taxonomy and a fabric taxonomy generated in collaboration with materials scientists. TextileNet can be used to train and evaluate the state-of-the-art Deep Learning models for textile materials. The goal of TextileNet is to contribute to textile material identification in the fashion industry and image-based textile material retrieval, at the same time, standardise the digitized textile material labeling. TextileNet can be deployed in various domains, including materials science, fashion design, retail and the textile supply chain, *etc.* TextileNet's contributions are:

- A fibre taxonomy and a fabric taxonomy created in collaboration with material domain experts; these taxonomies contain macro-types of textiles and are extendable for future new fibre/fabric types;
- Using the labels from these taxonomies, we collect and build material taxonomy-based fashion datasets for fibre and fabric. The built datasets, named TextileNet, TextileNet-fibre contains 33 fibre labels, 27 fabric labels in TextileNet-fabric, and have 760,949 images;



**Figure 2: (a) Traditional and (b) LLM-based conversational agents with voice inputs and outputs. The traditional agent has various components such as NLU, NLG and Dialogue database searching. In contrast, the LLM-based agent simply uses the database to generate prompts for the LLM, enabling a much simpler and easy-to-develop pipeline.**

- TextileNet presents and reports four baseline models (CNNs and Vision Transformers) for fibre and fabric classification, both models achieve > 87% top-5 accuracy on our datasets.

With the designed fibre and fabric taxonomies, one can now trace the origin and production flow of a particular fabric. To further illustrate how the fabric taxonomy and TextileNet-fabric could work in conjunction with the fibre taxonomy, in Figure 1, we show how a practitioner may use a trained ML model (TextileNet) to predict the fabric of garments, for instance, in this case, these are *flannel* and *lawn*. This information allows users to map the manufacturing flow (fabric taxonomy in green dash, fibre taxonomy in red line), showing how *lawn* connects to *nylon* and *flax*, which further connect to raw materials like *flax* and *petroleum*. This trace-back process not only characterizes the manufacturing process but also identifies raw material categories (*animal-based* or *synthetic*). These taxonomies help researchers and designers estimate the life-cycle and carbon footprints of garments. For example, carbon emission numbers or durability measurements can be assigned to transitions between fibre and fabric, such as *lawn* → *flax* and *lawn* → *nylon*. The carbon footprint can also be estimated based on these raw materials.

These taxonomies then help materialise the idea of creating a textile circular economy discussed in material science [24]. Future applications for this dataset range from textile classification to optimization of the textile supply chain and interactive design for consumers. We envision that this can contribute to the development of a new AI-based fashion platform.

#### 4 Let AI talk about textiles circularity – TextileBot


Developing domain-specific conversational agents has long been considered a challenge due to the need for large amounts of domain-focused data. Previous voice-based conversational agents have typically been built using Natural Language Understanding (NLU) and Natural Language Generation (NLG) along with database-supported dialogue management [1, 19] (see Figure 2a). However, the development of conversational agents is often hindered by the scarcity of data [3, 13, 26]. While Chatbot (text-based) interfaces have been studied extensively, comparatively little research has been done on natural spoken dialogue (voice-based interaction).

Recent advancements in pre-trained Large Language Models (LLMs) such as GPTs have made it possible to generate human-like responses to complex queries, making them a viable option as a knowledge backbone for such agents reducing the need for data collection and training. LLMs can perform various downstream tasks in a zero-shot fashion (*i.e.*, without any training) through prompting. We propose combining the power of LLMs with our domain-specific, taxonomy-based prompting strategies to tackle data scarcity. This approach uses LLMs as the knowledge backbone, eliminating the need for a dialogue database.

We chose to build a conversational agent in the textile domain to demonstrate our methodology. It has diverse information and expertise from various areas, *e.g.* fashion, home textiles, supply chain, materials science, and manufacture *etc.* Moreover, as the second-largest polluter globally, the textile industry urgently needs sustainable practices [2, 31]. Understanding sustainability aspects, like recycling, is challenging. However, conversational agents are widely used in fashion retail [1, 5], offering an opportunity to integrate socially responsible behaviours (*e.g.*, sustainability communication) into business strategies [8]. This domain can benefit from the technology socially and economically.

In this work [37], we first present a method to incorporate structural knowledge into prompts and demonstrate how LLMs equipped with these carefully crafted prompts can build domain-specific conversational agents, particularly for the textiles domain. We proposed an optimization technique that enables LLMs to interact with users with varying breadth and freedom. We also introduced a prompting refinement technique to facilitate LLMs engaging with users in voice-based interactions, allowing the LLM to retain statefulness, enabling a more natural, “multi-round” (continuous) conversation.

To showcase our approach, we integrated a prompted LLM into a physical voicebot called TextileBot. It is a voice interface that conveys the topic of textiles circularity, developed on a Raspberry Pi (see Figure 3a), which is capable of engaging users in dynamic conversations. We evaluate TextileBot’s effectiveness through an in-person user study (see Figure 3), focusing on a human-in-the-loop strategy to foster human-centred AI design [20]. Success was measured by the participants’ ability to discern the unique CA characteristics crafted through our approach and the system’s proficiency in maintaining domain-specific, multi-turn conversations.



(a) The physical TextileBot interface.



(b) A participant interacts with the TextileBot.

**Figure 3: Left: TextileBot** - The physical agent interface is composed of a 3D printed box (6), a speaker (5), a microphone (4), and a button (3), all integrated into the Google AIY board (2) mounted on the Raspberry Pi 3 Model B (1) presented in (a). **Right: A participant interacting with the TextileBot used across all three voice-based agents (b).**

This evaluation integrates both quantitative data and qualitative insights, uncovering the complex dynamics of these human-agent interactions and exploring various facets of human behaviour, engagement, and responses.

This research addresses two timely questions: First, while LLMs as foundation models for general-purpose conversational agents have been studied [17, 21, 35], their application in creating domain-specific voice agents is less understood. Second, there is a limited understanding of how humans perceive and interact with domain-specific voice-based conversational agents powered by prompted LLMs. These knowledge gaps arise from factors such as the complexity of natural languages, technical challenges with text-to-speech integration, handling ambiguity and context in human speech, and the need for effective user experience design to facilitate seamless and natural user interactions with the system [4, 32]. In summary, the main contributions of this work are three-fold:

- We develop a novel three-phase method that enables LLMs to transform from task-agnostic to domain-specific, adapting different conversational styles and incorporating memory for continuous conversation. These phases can be used individually or in combination, depending on the task at hand.
- We integrate the method into a design of a domain-specific voice-based device, *i.e.*, TextileBot. This voice-based agent is custom-designed for the textile circularity domain, facilitating tailored conversations with consumers regarding circular economy practices within the textile industry.
- We evaluated three variations of TextileBot via in-person interactions to assess our approach’s effectiveness and to understand user interactions. We analyzed these human-agent interactions, providing qualitative descriptions and quantitative insights. This analysis aims to inform potential design improvements in the broader domain of AI-enabled voice interfaces.

This tool engages users in conversations about textiles, offering a novel method for prototyping domain-focused conversational

agents. We hope our approach can be applied to other domains lacking extensive data.

## 5 Conclusion and next steps

Through this multifaceted research journey, I am privileged to be part of the multidisciplinary team at the UKRI Interdisciplinary Circular Economy Centre for Textiles. This doctoral research focuses on developing new technologies for digital experiences and enhancing Human-AI interactions. By bridging the gap between material science and computer science, the research aims to innovate machine capabilities, and raise human awareness through digital multisensory experiences.

Currently, I am midway through my third year. I am focusing on how AI perceives the physical world, including touch [36] and olfactory experiences. The next phase of my work involves enhancing AI’s integration with these multisensory experiences to facilitate more natural interactions. This can involve leveraging immersive technologies like virtual reality (VR) and augmented reality (AR) to create immersive natural multisensory environments. Specifically, the research will address the following questions: How does existing AI perceive and interpret sensory experiences, and how can we craft multisensory textile experiences that accurately map textile materials in sensory experiences with AI?

To achieve these goals, I will conduct comprehensive studies to assess how current AI models interpret and classify different textile materials. I will develop and test methodologies that incorporate multisensory inputs—such as touch, visual, and olfactory stimuli—into AI models to enhance their understanding of textile materials and properties. Additionally, I will design multisensory experiences that leverage human expertise to create more intuitive and natural AI interactions.

By addressing these steps, my PhD research aims to advance the state of AI in understanding and interacting with domain-specific practice, in my case, textiles, ultimately supporting the development of a sustainable fashion industry through innovative digital multisensory experiences.



## Acknowledgments

Thank you to my PhD supervisors, Prof Marianna Obrist, Dr Elia Gatti as well as Prof Youngjun Cho. This doctoral research is supported as part of the Engineering and Physical Sciences Research Council (EP/V011766/1) National Interdisciplinary Circular Economy Research (NICER) programme for the UKRI Interdisciplinary Circular Economy Centre for Textiles: Circular Bioeconomy for Textile Materials.

## References

- [1] Merav Allouch, Amos Azaria, and Rina Azoulay. 2021. Conversational agents: Goals, technologies, vision and challenges. *Sensors* 21, 24 (2021), 8448.
- [2] Kerrice Bailey, Aman Basu, and Sapna Sharma. 2022. The Environmental Impacts of Fast Fashion on Water Quality: A Systematic Review. *Water* 14, 7 (2022), 1073.
- [3] Ms Aayushi Bansal, Dr Rewa Sharma, and Dr Mamta Kathuria. 2022. A systematic review on data scarcity problem in deep learning: solution and applications. *ACM Computing Surveys (CSUR)* 54, 10s (2022), 1–29.
- [4] Amanda Baughan, Xuezhi Wang, Ariel Liu, Allison Mercurio, Jilin Chen, and Xiao Ma. 2023. A Mixed-Methods Approach to Understanding User Trust after Voice Assistant Failures. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [5] Rodrigo Bavaresco, Diógenes Silveira, Eduardo Reis, Jorge Barbosa, Rodrigo Righi, Cristiano Costa, Rodolfo Antunes, Marcio Gomes, Clauter Gatti, Mariangela Vanzin, et al. 2020. Conversational agents in business: A systematic literature review and future research directions. *Computer Science Review* 36 (2020), 100239.
- [6] Wen-Huang Cheng, Sijie Song, Chieh-Yun Chen, Shintami Chusnul Hidayati, and Jiaying Liu. 2020. Fashion Meets Computer Vision: A Survey. *arXiv preprint arXiv:2003.13988* (2020).
- [7] Desiree Cho, Shu Zhong, Madhan Kumar Vasudevan, and Marianna Obrist. 2024. Mindful Touch: Mid-Air Haptics Facilitates Novices' Experiences of Audio-Guided Mindfulness Meditation. Available at SSRN 4768789 (2024).
- [8] Mariachiarra Colucci, Annamaria Tuan, and Marco Visentin. 2020. An empirical investigation of the drivers of CSR talk and walk in the fashion industry. *Journal of Cleaner Production* 248 (2020), 119200.
- [9] Kirsti Cura, Niko Rintala, Taina Kamppuri, Eetta Saarimäki, and Pirjo Heikkilä. 2021. Textile recognition and sorting for recycling at an automated line using near infrared spectroscopy. *Recycling* 6, 1 (2021), 11.
- [10] Textile Exchange. 2021. *Preferred Fiber & Materials Market Report 2021*. Textile-Exchange.
- [11] The Ellen MacArthur Foundation. [n. d.]. Fashion and a circular economy | Ellen MacArthur Foundation. <https://ellenmacarthurfoundation.org/topics/fashion/overview>
- [12] Nancy Frey and Douglas Fisher. 2008. *Teaching visual literacy: Using comic books, graphic novels, anime, cartoons, and more to develop comprehension and thinking skills*. Corwin Press.
- [13] Alexander Frummet, David Elsweiler, and Bernd Ludwig. 2022. "What Can I Cook with these Ingredients?"—Understanding Cooking-Related Information Needs in Conversational Search. *ACM Transactions on Information Systems (TOIS)* 40, 4 (2022), 1–32.
- [14] Sheng Guo, Weilin Huang, Xiao Zhang, Prasanna Srihanta, Yin Cui, Yuan Li, Hartwig Adam, Matthew R Scott, and Serge Belongie. 2019. The imaterialist fashion attribute dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*. 0–0.
- [15] Fu Jia, Shiyuan Yin, Lujie Chen, and Xiaowei Chen. 2020. The circular economy in the textile and apparel industry: A systematic literature review. *Journal of Cleaner Production* 259 (2020), 120728.
- [16] Jon Kabat-Zinn. 1982. An outpatient program in behavioral medicine for chronic pain patients based on the practice of mindfulness meditation: Theoretical considerations and preliminary results. *General hospital psychiatry* 4, 1 (1982), 33–47.
- [17] Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert McHardy. 2023. Challenges and Applications of Large Language Models. *arXiv preprint arXiv:2307.10169* (2023).
- [18] Asimananda Khandual and Swikruti Pradhan. 2019. Fashion brands and consumers approach towards sustainable fashion. In *Fast fashion, fashion brands and sustainable consumption*. Springer, 37–54.
- [19] Sheetal Kusal, Shruti Patil, Jyoti Choudrie, Ketan Kotecha, Sashikala Mishra, and Ajith Abraham. 2022. AI-based Conversational Agents: A Scoping Review from Technologies to Future Directions. *IEEE Access* (2022).
- [20] Mina Lee, Megha Srivastava, Amelia Hardy, John Thickstun, Esin Durmus, Ashwin Paranjape, Ines Gerard-Ursin, Xiang Lisa Li, Faisal Ladhak, Frieda Rong, et al. 2022. Evaluating Human-Language Model Interaction. *arXiv preprint arXiv:2212.09746* (2022).
- [21] Peter Lee, Sebastien Bubeck, and Joseph Petro. 2023. Benefits, limits, and risks of GPT-4 as an AI chatbot for medicine. *New England Journal of Medicine* 388, 13 (2023), 1233–1239.
- [22] Jared R Lindahl, Nathan E Fisher, David J Cooper, Rochelle K Rosen, and Willoughby B Britton. 2017. The varieties of contemplative experience: A mixed-methods study of meditation-related challenges in Western Buddhists. *PLoS one* 12, 5 (2017), e0176239.
- [23] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. 2016. Deep-fashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1096–1104.
- [24] Ellen MacArthur et al. 2013. Towards the circular economy. *Journal of Industrial Ecology* 2 (2013), 23–44.
- [25] Ellen MacArthur et al. 2014. Towards the circular economy: Accelerating the scale-up across global supply chains. In *World Economic Forum*.
- [26] Selina Meyer, David Elsweiler, Bernd Ludwig, Marcos Fernandez-Pichel, and David E Losada. 2022. Do We Still Need Human Assessors? Prompt-Based GPT-3 User Simulation in Conversational AI. In *Proceedings of the 4th Conference on Conversational User Interfaces*. 1–6.
- [27] Subramanian Senthilkannan Muthu. 2018. *Circular Economy in Textiles and Apparel: Processing, Manufacturing, and Design*. Woodhead Publishing.
- [28] Kirsi Niinimäki, Greg Peters, Helena Dahlbo, Patsy Perry, Timo Rissanen, and Alison Gwilt. 2020. The environmental price of fast fashion. *Nature Reviews Earth & Environment* 1, 4 (2020), 189–200.
- [29] Donald A Norman and Stephen W Draper. 1986. User centered system design: New perspectives on human-computer interaction. (1986).
- [30] Nynne Norup, Kaj Pihl, Anders Damgaard, and Charlotte Scheutz. 2019. Evaluation of a European textile sorting centre: Material flow analysis and life cycle inventory. *Resources, Conservation and Recycling* 143 (2019), 310–319.
- [31] AK Roy Choudhury. 2014. Environmental impacts of the textile industry and its assessment through life cycle assessment. In *Roadmap to sustainable textiles and clothing*. Springer, 1–39.
- [32] Katie Seaborn, Norihisa P Miyake, Peter Pennefather, and Mihoko Otake-Matsuura. 2021. Voice in human-agent interaction: A survey. *ACM Computing Surveys (CSUR)* 54, 4 (2021), 1–43.
- [33] Madhan Kumar Vasudevan, Shu Zhong, Jan Kučera, Desiree Cho, and Marianna Obrist. 2023. MindTouch: Effect of Mindfulness Meditation on Mid-Air Tactile Perception. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [34] Carlos Velasco and Marianna Obrist. 2020. *Multisensory Experiences: Where the senses meet technology*. Oxford University Press.
- [35] JD Zamfirescu-Pereira, Heather Wei, Amy Xiao, Kitty Gu, Grace Jung, Matthew G Lee, Bjoern Hartmann, and Qian Yang. 2023. Herding AI Cats: Lessons from Designing a Chatbot by Prompting GPT-3. (2023).
- [36] Shu Zhong, Elia Gatti, Youngjun Cho, and Marianna Obrist. 2024. Exploring Human-AI Perception Alignment in Sensory Experiences: Do LLMs Understand Textile Hand? *arXiv preprint arXiv:2406.06587* (2024).
- [37] Shu Zhong, Elia Gatti, James Hardwick, Miriam Ribul, Youngjun Cho, and Marianna Obrist. 2024. LLM-Mediated Domain-Specific Voice Agents: The Case of TextileBot. *arXiv preprint arXiv:2406.10590* (2024).
- [38] Shu Zhong, Miriam Ribul, Youngjun Cho, and Marianna Obrist. 2023. TextileNet: A Material Taxonomy-based Fashion Textile Dataset. *arXiv preprint arXiv:2301.06160* (2023).
- [39] Xingxing Zou, Xiangheng Kong, Waikung Wong, Congde Wang, Yuguang Liu, and Yang Cao. 2019. Fashionai: A hierarchical dataset for fashion understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 0–0.