Drinking Chai with Your (AI) Programming Partner: Value Tensions in the Tokenization of Future Human-AI Collaborative Work

Michael Muller
IBM Research AI
Cambridge, MA, USA
michael_muller@us.ibm.com

Justin D. Weisz
IBM Research AI
Yorktown Heights, NY, USA
jweisz@us.ibm.com

Stephanie Houde
IBM Research AI
Cambridge, MA, USA
Stephanie.Houde@ibm.com

Steven I. Ross
IBM Research AI
Cambridge, MA, USA
steven_ross@us.ibm.com

ABSTRACT

The new technologies of generative AI present important ethical challenges for application design in relation to human work and future workplaces. Using an extended design fiction, we conducted a critical inquiry into possible futures of generative applications based in human activities. We explored potential benefits and harms. 30 members of our international workplace read the design fiction, rated each of the 13 brief chapters, and critiqued the narrative. We interpret their responses using value tensions (Value Sensitive Design) and design theories of refusal, non-use, and un-use. We propose *situated refusal*, a combination of value tensions and refusal theory, and we describe three major dimensions of that perspective: Who is guiding the work? How much voluntary adoption? Who else is affected by the individual worker's decisions?

CCS CONCEPTS

• **Human-centered computing → HCI design and evaluation methods**; *Human computer interaction (HCI)*; **Interaction design process and methods**; • **Computing methodologies → Artificial intelligence**; *Artificial intelligence*; • **Software and its engineering → Software development techniques**; **Software development techniques**.

KEYWORDS

Design fiction, Generative AI, Future of work, Workplace, Labor, Displacement, Value sensitive design (VSD)

ACM Reference Format:
Michael Muller, Justin D. Weisz, Stephanie Houde, and Steven I. Ross. 2024. Drinking Chai with Your (AI) Programming Partner: Value Tensions in the Tokenization of Future Human-AI Collaborative Work. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work (CHIWORK '24), June 25–27, 2024, Newcastle upon Tyne, United Kingdom.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3663384.3663390



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

CHIWORK '24, June 25–27, 2024, Newcastle upon Tyne, United Kingdom
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1017-9/24/06
https://doi.org/10.1145/3663384.3663390

1 INTRODUCTION

> "Might we allow the [AI technology] to be not just good or bad, submissive or dominant, but a complicated, contradictory being?... The future world serves as a projection of current issues and conflicts, and thus the future becomes a way of looking at ourselves and our culture."
> – Søndergaard and Hansen [130]

There are many work domains in which people may use emerging AI technologies, or people may be used by those technologies. As Butler et al. [17, p. 40] recently wrote, "Instead of 'how *will* AI affect work?', the question should be 'How *do we want* AI to affect work?'" (italics in the original text). This question of human agency is increasingly important for the future of work and workplaces, where AI-based technologies are already disrupting ways-of-working and human lives [16, 17, 51, 64, 116, 125]. Since the era of expert systems [6, 11, 117, 118], there has been a risk that human skills may be appropriated into technology, deskilling the workforce [16, 81, 113, 144]. The risk is increased with more recent AI technologies that can "learn" human work through observation and modeling [1, 66, 69]. This removal of knowledge from workers into technology, has the effect of appropriating job skills and work knowledge from labor, thereby strengthening management's power in job design and labor-management relations [32, 102, 108, 119].

A recent technological development may accelerate these changes. Creators of generative AI use deep neural networks to compute models of complex data domains, including human work domains [20, 36, 58, 118, 125]. Generative AI technologies have matured at a rapid pace, allowing humans to create high-fidelity content from simple natural language descriptions [90, 106, 109]. In other domains, high-quality models have been developed that help people write textual works [18, 43], create musical compositions [45, 84, 85, 123], create and refine images [71, 82, 101] and videos [5]. Domain-specialized models can help identify promising molecular compounds for developing new drugs [22, 29, 89], and even write and document software programs [110, 143]. And yet, the question

remains unanswered: What is the nature of the technology, and who or what is in control?

Thereby, these powerful technologies present us with renewed ethical issues, including questions of ownership of employees' knowledge [41, 59, 125] and employment risks such as de-skilling and job elimination [21, 62, 72, 113] through dynamics of data colonialism [96, 138, 142]. In view of the positive light in which generative AI technologies are cast in popular press (e.g. [46, 47, 92, 103]), we are interested in pursuing a more critical view of the potential benefits and harms of future generative applications.

In this paper, our goal is to explore how maturing generative AI technologies may impact the future human experience and sustainability of work. To pursue this goal, we use an extended design fiction (e.g., [86, 133]) to explore possible futures of generative AI in software workplaces, and by implication other workplaces as well. We report assessments and critiques of the fiction by an international sample of colleagues whose work could be affected by these technology developments, interpreted through Value Sensitive Design (VSD) [37, 38] and specifically the Value Tensions concept [91] of VSD. We based the fiction on our own experiences in working for an international technology company, as well as new software engineering and data science tools (e.g. [65, 68, 135]) that have the potential either to serve as a valuable assistant, or in the extreme case, to replace the workers whom they were designed to assist.

We examine these conundrums as instances of what Miller et al. described as *value tensions* [91] (an aspect of value sensitive design [38]). Through that theoretical lens, we analyze the employee-institutional relation in terms of convergent interests (value flows) and divergent interests (value dams). In some value dams, employees may wish to reduce or refuse certain aspects of these new technologies. We apply the research traditions of refusal [41], non-use [104, 115], and un-use [75] for insights into these choices. However, we also note that nearly all of the refusal literature deals with consumer choices and voluntary use. Therefore, we explore limitations on employees' ability to negotiate their refusal or resistance certain institutionally-designated technologies [79, 83]. Based in values tensions, we propose a reconsideration of the concepts of refusal, non-use, and un-use – namely, situated refusal.

Our paper makes the following contributions to the HCI community:

- an evolving set of speculations regarding future generative AI applications in software engineering workplaces, structured as a design fiction about the experiences of a worker in such a future workplace;
- an empirical assessment of these speculations from the perspective of potential future workers;
- a proposed extension to theories of value tensions and refusal, situated refusal.

2 BACKGROUND

We describe the three domains that we bring together in this paper: *generative AI* and *refusal*, explored through *design fiction*. We tie them together through concepts of value tensions and standpoint perspectives from feminist science and technology studies.

2.1 Generative AI and Tokenization

Generative AI is a recent computing method that focuses on the creation of novel artifacts as the outcome. By contrast, discriminative AI focuses on determining a decision boundary between one or more classes or outcomes (e.g. is an image of a hot dog, or should a mortgage application be approved). Generative models are usually trained using a large number of examples from a domain [44], which are broken down into a sequence of components or elemental *tokens* - i.e., the components are tokenized. These tokens are then encoded as high-dimensional vectors in a *representation space* (or *latent space*).

The next, crucial step is that each token is predicted (i.e., modeled) on the basis of the preceding tokens. After that step, then the model can predict a word, design element, musical phrase, etc., on the basis of preceding or surrounding words, design elements, phrases, etc. Researchers and developers have used generative technologies to create designs [7, 146], images [71, 82, 101], music [45, 85, 123], texts [18, 43, 143], videos [5], source code [110], and molecular descriptions [29, 89]. In general, these approaches have involved the strategy of predicting the next token (i.e., pixel, word, code-fragment, molecular structure) on the basis of previous such tokens. Once a domain has been *learned* in this way, then specific instances or "solutions" may often be implemented through a few additional data records - a process called "few-shot" or "one-shot" learning [66, 69].

In this paper, we consider how this approach might be applied not only to "data" (words, images, pixels...) but also to human activities. We use concepts from value sensitive design to clarify some of these issues.

2.2 Value Tensions

Focusing on the diversity among human values, we speculate on near-term future workplaces in which human actions may be similarly tokenized and learned by generative applications, supporting subsequent specialization for particular tasks or patterns of work through few-shot or one-shot learning approaches. We shift the discussion from "instances of *what*" to the more critical question, "instances of *whom*?"

Through thirteen chapters in a design fiction, we consider miniscenarios that might perceived differently by workers and by management. To analyze these potential disagreements, we make use of Miller et al.'s concept of *value tensions* [91]. Miller et al. work in the theoretical framework of Value Sensitive Design (VSD) [38], which encourages an analysis of the values brought by diverse *stakeholders* who are directly or indirectly involved in a technology or a policy. VSD can help us to understand the values that each stakeholder brings to a potentially contested situation. Value tensions focus on those contestations. In Miller et al.'s analysis, a *value flow* occurs when there is substantial agreement about the values relating to a technology or a policy. By contrast, a *value dam* occurs when there is substantial disagreement.

Value tensions have been used to structure analyses of existing projects, such as production line design [74], workload distribution among teams [53], sensitive personal data in healthcare [4, 126] and parenting [28], and urban design [132]. In addition to present-day dilemmas, Nathan et al. applied value tensions to examine

possible futures for designs and technologies [99] using a fictive approach called *value scenarios* [100]. While most studies have examined value tensions *between* stakeholders, Smits et al. [126] and Friedman et al. [37] note that individuals can experience *internal* value tensions between contrasting goals or priorities. In this paper, we apply value tensions as an analytic tool to examine potential future uses of a powerful technology (similarly to [99]), and we consider how some value conflicts may occur within an individual worker's decisions (similarly to [126]).

2.3 Refusal, Non-Use, Un-Use

In situations with substantial value dams, workers may want to refuse or resist technologies or policies. Refusal may be considered as a simple negation or foreclosure of action (e.g., [25, 41, 104]). However, scholars have also argued that refusal can be a creative or even a generative act [10, 42, 87, 140]. Indigenous refusal is, according to Tuck and Yang, a generative stance in response to forms of colonial oppression [141]. Sabie et al. refer to unmaking (a form of refusal by disassembly or destruction) as a kind of material innovation, including uncrafting and unfabricating [112]. In the related concept of un-use, Krischkowski et al. reflect on "activities of un-use as a kind of making" - i.e., a making of alternatives to no-longer-desired actions of use [75].

In the CHI and related literatures, refusal can take many forms:

- In **Data Refusals**, individuals find ways to avoid contributing their data - e.g., a refusal to use certain websites, or to allow personal data to be commercialized [25].
- **Knowledge refusals** have been discussed in HCI at least as long ago as the Collective Resource approach in participatory design, in which labor withheld shopfloor knowledge from management until management were willing to engage in democratic decision-making about jobs and work-design [32]. Once management have obtained the knowledge of the workforce, they may inscribe or operationalize that knowledge into technologies, effectively transferring informed control of work from the workforce and into management's algorithms [49, 50, 70, 72, 97]. Knowledge refusals are also an important aspect of the broader actions in Indigenous refusal, as a resistance to colonial exploitation of Indigenous and Aboriginal knowledges [24, 124, 141].
- **Designer/developer Refusals**, in which "the implication is not to design" [10], have been advocated as an ethical stance in relation to technologies that should not be built or deployed.
- In **Disenchantment**, workers prefer an earlier way of doing their jobs [75, 115].
- Workers may also become **Disinterested** from using a specific technology [75, 78, 96, 115].

Multiple forms of refusal and non-use take place in this paper's design fiction, and are discussed by readers of that fiction.

2.4 Design Fictions

We use design fiction as our method of inquiry to ask whether organizations or societies may wish to place limits on AI-based generation of certain types of artifacts or actions [3, 27, 41, 105].

For more than a decade, scholars in HCI have used design fictions as a core set of research methods for critical exploration of sociotechnical choices and their consequences [39, 86, 133], often taking the form of extended "what-if" speculations [73, 86], including speculations that warn of unwanted outcomes [127]. Work in design fictions has matured to the point that there are now discussions of protocols to evaluate them [8, 94]. Skirpan and Fiesler [121] used design fictions to encourage researchers to think about possible extractive and manipulative applications of digital sensing technologies. Houde et al. [60] considered ways in which generative technologies might be used to violate norms of authorship and evidence, and to solicit colleagues' views on those questions. Søndergaard and colleagues addressed questions of gendered technologies, including issues of *who* designs and deploys those technologies [128–130]. Muller and Liao [95] used a participatory approach to learn about the values that motivated how participants wrote the conclusions of half-written stories abouut futuristic technologies and usages. Inverting this process, Cheon and Su [23] asked participants to speculate about their imagined histories of a future situation could have come to be. Jensen and Vistisen [63] took an explicitly ethical perspective. Stories invite active, engaged thinking about ambiguous situations and ethical challenges.

In this paper, we focus on particular technologies and their implications for work. We orient our explorations through the focus on persons, their perspectives, their needs, and their contestations. Based in earlier discussions of how an individual's socioeconomic position can influence perceptions and actions [33, 88], Feminist standpoint theory makes a strong case to consider that perspectives and even "facts" are based on *who* is taking each perspective and correspondingly, *who* is determining each "fact" [52, 55]. These disagreements may lead to value dams in which two or more parties disagree about the ethics of a particular technology.

Working in the VSD framework, Nathan et al. used *values scenarios* to argue that work with designed technologies does not always turn out as expected by the makers of those technologies [100], with very different outcomes for different stakeholders. Fanni et al. described related tensions between *active* human agency (i.e., human action to monitor and correct an AI's actions) vs. *passive* human agency (i.e., deferring control to the AI) [34]. Value dams and value flows [91] help us to consider how the values of each stakeholder constituency (i.e. employees and organizations) are based in part on their perspective on the work and the workplace; for example, an organization's desire for control and predictability vs. an employee's desire for autonomy and the protection of the skilled basis of their work [113, 147]. Although the principles of generative AI may or may not be benign, core ethical issues arise in human and societal choices in how to apply those principles [2, 60]. We utilize the design fiction methodology to help us consider some of those choices.

3 DRINKING CHAI WITH YOUR (AI) PROGRAMMING PARTNER

We wrote this fiction based on an on-going project that explores multiple workplace applications and future implications of generative AI. One author wrote the original fiction. Three authors

commented on drafts, leading to major changes. Some changes improved the focus on the technology, while other changes improved the narrative flow.

We now present our design fiction. For ease of participants' commentary in our empirical examination of the fiction (Section 4.2), we divided the narrative into thirteen brief chapters.

The rationales for the components of the fiction are as follows.

- **Chapter 1** introduces the situation and the principal human and AI actors.
- **Chapters 2**, **4**, and **5** explore how generative AI can assist with a typical task for software engineers.
- **Chapter 3** introduces the theme of differing strategies among human and AI; this theme is pursued further in **Chapter 8**.
- **Chapter 6** further explores how task-tokenization via generative AI can support a human, through a second typical task in the work of software engineers.
- **Chapters 7** and **9** introduce issues around intrusive "contributions" by the AI.
- **Chapters 11** and **12** introduce the theme of displacement, in which a tokenized model of Tikaani's work could replace Tikaani.
- **Chapter 13** concludes the trend toward dystopic interventions.
- **Chapters 3** and **7-13** explore various conditions and limited options of situated refusal.

Chapter 1: Introducing Tikaani and Your Programming Partner
Tikaani had been postponing work to translate the UX of the enormous BiggerFin application that their team was modernizing from its legacy Cobol code. So much had changed in the decades since BiggerFin was originally deployed.

Everyone who had been on the original BiggerFin team had retired or moved on to other jobs. Generative translation technologies had become increasing accurate to convert the back-end code, but modernizing the UI was still a challenge for AIs.

The UI remained a challenge because web-based UI technologies were now much better, and corporations' expectations for what was considered "good design" had also changed. Today was the day to do this! And the task would distract Tikaani from the loneliness of the office, with so many unoccupied desks. Fortunately, Tikaani had a conversational generative AI (genAI) assistant that could help, named "Your Programming Partner" (YPP).

Tikaani opened the spec from the design team, and called up YPP. YPP displayed "Greetings follow. Good morning, Tikaani. Today we will have a good day, a very good day," and followed with "Advertisement follows. Try BerryBytes in the ByteBar. They're brainfood. They can reduce any unwanted aftertastes."

Chapter 2: Analyzing the Spec
During the previous modernization project, Tikaani had told YPP that its nickname was going to be "Y". Tikaani pressed the speech-to-text key, and said "Y, please review the spec," and made a gesture to tell YPP which document was the spec.

"Drink chai," said YPP. Tikaani took a sip of chai.

After a brief pause, YPP displayed, "To perform a generative translation from spec to architecture. a set of examples is required. Recommendations follow: Finance/big-institution; Finance/small-institution; Corporate; SMB; or say 'other' for non-recommended domains." Each recommendation was displayed as a selectable button, so that Tikaani's choice would immediately lead to action by YPP.

Chapter 3: Tikaani's Social Search Strategy
While it was true that BiggerFin was indeed a large financial client and would have been a partial match to "Finance/big-institution," the clients of BiggerFin for this service would be people and small businesses seeking micro-loans. Tikaani wanted to use a more social search strategy to make their choice. There was no one nearby to ask, so Tikaani typed, "Display team members along with the number of financial projects they have worked on."

YPP produced the list, again making each row of the displayed table into a button-for-action. Tikaani saw that Yu had worked on 20 similar projects. Yu had been Tikaani's mentor when they joined the company, and Tikaani missed the days when they had traded ironic comments over tea.

Tikaani gestured to the "Yu – Finance projects" row, and YPP accepted this gesture as a choice+command, and began to work.

YPP displayed, "Further guidance is required. Recommendations follow. (1) Choose the granularity of microservices from this list: micro, mini, or macro. Macro is recommended. (2) Optimization can be a weighted sum of: performance, maintainability, alignment-to-code, alignment-to-usage/natural-seams, alignment-to-data. Balanced weights are recommended."

Tikaani revised the granularity to micro, and accepted YPP's recommendations about optimization weights.

Chapter 4: YPP Remembers Tikaani's Work Practices
YPP asked, "You usually prefer a JupyterGen notebook as your IDE [Integrated Development Environment]. Decision required: Should a JupyterGen notebook with your usual defaults be used for this project?"

Tikaani indicated "yes."

Chapter 5: YPP Anticipates Tikaani's Next Actions
YPP opened a JupyterGen notebook on Tikaani's display, and prepopulated the first cell with relevant Python libraries. Some of the libraries were open source, and some were proprietary libraries used to brand the modernized UXs as the company's products. YPP wrote major module names into the markdown

cells that preceded each code cell, and added draft documentation for the classes that each code cell would contain. YPP wrote stub code into some cells, and proposed full implementations into other cells.

"Drink chai," said YPP. Tikaani took a sip of chai.

Chapter 6: YPP Begins to Write Code

As YPP began to display the coding strategy into the shared Slack channel, Tikaani finished their sip of chai and put their cup down. Now that YPP was doing most of the work, Tikaani could relax and watch it unfold. They knew that they would need to review YPP's code, but they hadn't had to correct YPP's architectural assumptions during the past year. YPP called Tikaani's attention to lower-confidence code, and Tikaani made edits as needed. Tikaani was impressed that "Y" had learned so much of Tikaani's individual coding practices during their last year together. It was true that Tikaani missed the days when people actually had to solve their own coding problems without genAI support. But work with "Y" was easier, and much faster, and actually produced fewer bugs. This day of work with "Y", while maybe a little bit boring, was going to be a highly productive and also pleasurable – perhaps, indeed, a very good day.

YPP displayed, "Recognition follows. You earned 5 BerryByte points. Points were downloaded to your ID chip. Don't forget to redeem them at the ByteBar." And after a moment, "Message follows. CyberHR hope that this recognition will help you to have a good day, a very good day."

"Drink chai," YPP said. Tikaani took a sip of chai.

Chapter 7: YPP Proposes an Alternative Package

Tikaani began to perform a detailed review of the code that YPP had generated. As Tikaani opened a generic class, YPP displayed, "There are alternative modules to consider. Actionable Explanation follows. Class QuikClientPortfolio() from commercial package QuikFinance would be 20% faster, but has license fees that are usually $5k to $8k. Do you want to use QuikClientPortfolio() and submit a request for approval?"

Chapter 8: Tikaani Seeks Colleagues' Advice

Unsure of how to make the decision, Tikaani typed, "Inquire through YourPartnetNet for colleagues who have knowledge of QuikClientPortfolio. Poll 'Which release of QuikClientPortfolio is stable?' and indicate my status as Blocked. Use probable-knowldege feature. Anonymous responses are acceptable."

A minute later, YPP displayed, "6 colleagues responded. 5 said Release 15.3 is stable. 1 said Release 15.1 is stable. Two non-anonymized names are available."

Tikaani decided to accept the majority opinion. They typed, "Estimate likelihood of approval for QuikClientPortfolio Release 15.3, based on data from the current year only."

YPP displayed, "Estimation follows. Package QuikFinance has been approved on 82% of projects this year. Based on your personal history with requests, likelihood of approval for you is 90%." Then YPP repeated, "Do you want to use QuikClientPortfolio() and submit a request for approval?"

Tikaani pressed the speech-to-text key and said, "Paste QuikClientPortfolio, generate documentation of the decision rationale via GPT-neo and display it." YPP generated the rationale. Tikaani reviewed, and made two corrections. Tikaani said "Send request for approval, and add to my personal log." YPP displayed each step as it completed it.

"Drink chai," YPP said. Tikaani sipped.

Chapter 9: YPP Enforces Corporate Policy

Tikaani found a serious domain-related problem in one of YPP's generated modules. Apparently YPP didn't understand that micro-financing often involved clients with little collateral and incomplete credit histories. Tikaani knew that the loan-approval model would need to include additional factors. They began to rewrite the needed class from scratch. YPP issued a first warning: "The recommended module for this functionality is based on the MonthProjection() class from the FinBlast library." Tikaani continued to write code.

YPP continued, "Explanation follows. The estimated cost in work time for writing your own version is 3-5 hours for you, plus 2-4 hours for the QA team to test your new code, with a heightened risk of bugs of 34%. Required action follows. (1) provide rationale for writing your own version of this module; then (2) shift your work to a different module while (3) your rationale is automatically sent to your team-lead for approval. Full disclosure follows. There is a possibility that your team-lead will need to (4) auto-escalate your rationale to higher human management for further review. This may include a review of your user profile, with possible modification to your profile."

Chapter 10: Tikaani Responds

Tikaani sighed. They wrote the rationale, and moved on to the next one of YPP's generated documentation markdown cells. This day was looking less pleasurable.

Chapter 11: Virtual Employees

YPP played an audio clip and displayed, "CyberHR offer follows. You could be eligible for the BOGIE program. Buy Out Generative Image of Employee can enroll you at 150% of your base pay to help create a digital employee with your skills – a virtual you! Your work records are applied automatically as training data, and you may be asked to fine-tune the model.

You receive 150% of your base pay for the 12 months of development, which you receive as your buy out bonus when your employment terminates at the end of those 12 months. You agree that the contents of the model become the intellectual property of the company. Depending on how well your BOGIE image performs, you may be asked to stay on to continue the fine-tuning as an external consultant. Contact CyberHR to find out if you qualify for this exciting program."

The display was in a modal dialogbox, so Tikaani had to reply "I'm interested" or "Ask me later" to return to their work.

Chapter 12: Tikaani Declines the Buy-Out Offer
After the frustration of having to justify a necessary rewrite, Tikaani was almost ready to take the buy-out offer. But no. When the work was good, it was still interesting – even though it was lonelier now that Yu had taken the buy-out and had been replaced with v-Yu. Unlike Yu, v-Yu had no sense of humor or irony, and tended to respond to new situations with old strategies. v-Yu might be a co-*worker*, but was not the friend that Tikaani remembered with a sense of loss. Tikaani declined the offer, and went back to coding.

"Drink chai," YPP said. Tikaani glared at the screen.

Chapter 13: YPP Reminds Tikaani about a Workplace Wellness Requirement

"Drink chai," YPP repeated, "You are at 46%. Explanation follows. Having a good day is an employee responsibility, to maintain employee engagement for all. Chai helps. Below 40%, level-1 reporting to management is mandated. Required action follows. Drink chai." Was that a robot joke? Usually YPP accessed an external jokes database and announced "joke follows," but maybe the engineers were adding humor to the core conversational features? Tikaani had heard rumors of an experimental emotion module. Just before Yu had retired and been replaced with v-Yu, Yu had suggested that —

YPP interrupted Tikaani's reverie. "Drink chai. You are at 42%."

Tikaani gulped down a full mouthful of chai, and reached for a BerryByte to reduce the metallic aftertaste. Tikaani thought back to the days when chai was just an optional beverage that they drank with Yu's team during tea-breaks. Now chai had become a mandatory delivery vehicle for the Computer-Human Adaptive Intracellular. The short-lived CHAI virobots were linked to the cyberHR department through near-field communications to Tikaani's laptop. The virobots surveilled and corrected employees' health and psychological engagement, based on a generated target bio-labor profile from a learned dataset of biosignals from employees with demonstrated productivity and sufficient job-tenure.

Tikaani prepared for the brief moment of dizziness as the virobots crossed the blood-brain barrier to reach the emotion centers of the limbic system. YPP's warnings had left Tikaani in a very grim mood. But now, Tikaani felt the Intracellular beginning to activate. Tikaani experienced a familiar calm. Despite the loneliness of the mostly-empty office, it was going to be a good day after all, a very good day.

4 EMPIRICAL EVALUATION AND CRITIQUE OF THE DESIGN FICTION

The design fiction satisfied our desire to explore complex value tensions and ethical issues. However, we were uncertain about how others would interpret and reflect upon the future we constructed. Thus, we collected survey-based feedback on each story chapter to collect divergent viewpoints.

4.1 Method

4.1.1 Participants. We recruited a diverse sample of participants from within our organization by advertising our research in four company-wide Slack channels targeted at software engineers, designers, multicultural women, and researchers. We offered the equivalent of US $7.50 for less than 30 minutes of work (effective pay rate of US $15.00 per hour).

Thirty of our colleagues voluntarily agreed to participate. Participants who self-selected into our sample worked in 11 countries on four continents. 21 participants reported their preferred pronouns as she/her, and 9 reported as he/him. Participants worked in diverse software engineering roles[1] (50%) and software design (33%), with some in software research (13%) and other corporate functions (4%). 23 participants worked in the role of an individual contributor, 5 worked in the role of a manager, and 2 worked in the role of an executive.

4.1.2 Survey. The survey began with a page of demographic questions, and then displayed each of the 13 chapters on a separate survey page. For each chapter, we asked six Likert questions, which we analyzed in three pairs of ratings. Table 1 shows the six Likert questions. The first pair of questions (scenario and features) was intended to analyze the narrative flow of the fiction vs. the imagined new technologies. The second pair of questions (Tikaani vs. you) was intended to analyze the Tikaani character in situ vs. the readers' own, possibly different responses had *they* been in Tikaani's situation. The third set of questions (company vs. employee) directly addressed the value tensions in the story, within the broader framework of Value Sensitive Design.

To invite participants to think critically about the content of each chapter - or to protest against the unfolding story - we followed the six questions with "Do you think this part of the story should be different? If yes, please tell us what the difference(s) should be."

In summary, each of the 13 chapters was followed by six Likert-scale questions and one free-text discussion question.

4.2 Analysis

4.2.1 Likert Scales. Details of the quantitative analysis of the Likert-scale questions are shown in Figure 1. We analyze the results in three pairwise subsets of the data, in experimental designs of

[1] In this company, the work of software engineers is quite broad and includes activities such as programming, testing, and designing software architectures.

- 2 topics
- 13 chapters
- topics *x* chapters

We treat the 13 chapters as a narrative storytelling arc, and we test linear and quadratic trends across the 13 chapters. We also examine the interaction of topics *x* chapters, to test if there are differences in the trends for each of the two topics in each analysis.

We conducted three separate analyses of variance (ANOVA) - one for each of the pairs of questions in Table 1. Each pair answers a particular contrasting question. The first pair of Likert-scale questions contrasts the likelihood of the storyline with the likelihood of the technology features. The second pair of Likert-scale questions contrasts the actions of the protagonist, Tikaani, with what the participant thought that they themself would do. The third pair of Likert-scale questions contrasts the benefit to the employee (Tikaani) vs. the benefit to the company.

Participants gave significantly different responses across the different chapters. In general, participants' scores decreased from chapter 1 to chapter 13, indicating less credibility and less approval for the later parts of the narrative. For all three pairs of questions, the linear decreasing trend was significant, and the quadratic trend was also significant, indicating that the decrease became steeper toward the later chapters.

The ANOVAs allowed us to test for differences between the categories within each pair of questions. Participants responded that they themselves would be significantly less likely to act the way that Tikaani did. Participants also responded that, across the 13 chapters, the company was more likely to benefit than the employee was.

There was a small significant interaction effect of Chapters *x* the pair of (Scenarios vs. Features), mostly attributable to the lower likelihood of the features in chapter 12 ("Tikaani declines the buy-out offer"). There was also a small significant interaction effect for chapters *x* Company vs. Employee, mostly attributable to a sharp contrast at chapter 11, where participants report that the Company was much more likely than the Employee to benefit from the virtual employee program.

In general, participants found the more futuristic chapters to be less plausible and less beneficial. We now turn to the participants' comments on each chapter, for a richer interpretation.

4.2.2 Participants' Comments. To analyze participants' comments, we followed the 2018 version of Reflexive Thematic Analysis (RTA) [15, 26] for the open-text responses for each of the thirteen chapters, supplemented by thematic mapping [137]. The open-texts were brief, often a sentence or less, such as, *"Yikes. Hopefully improbably"* (P12, Chapter 13). As summarized in the thematic map [14] of Figure 2, we initially identified six open codes. Through subsequent iterations, we split three of those six open codes into two subcodes (e.g., from "Plausibility" we constructed "Plausible advertisements" and "Plausible technologies"). One open code ("Readiness to use AI features") was used so infrequently that we removed it. In our final iteration, we constructed two themes, of "Human Control" and "Human Displacement." Figure 3 shows a mapping of subthemes to chapters. These mappings are for the strongest themes in each chapter. In a few cases, a theme was mentioned in a minor way, and this mapping does not indicate those infrequent mentions. (Note:

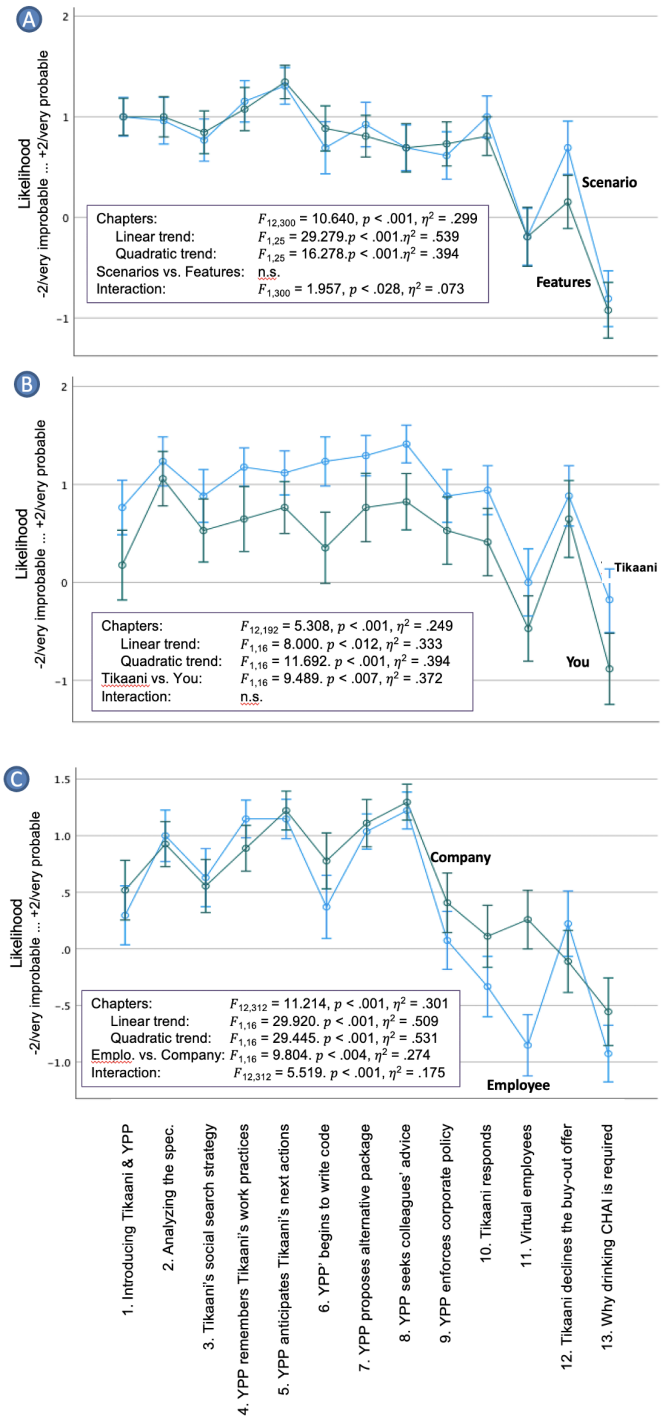


Figure 1: Ratings data from six questions, across the thirteen chapters. A. Likelihood ratings of Scenario and Technology Features. B. Likelihood ratings of Tikaani's actions and Your (Respondent's) actions. C. Likelihood ratings of benefit for Employees and Company.

Likelihood of **Scenario and Features**: - **[Scenario]** Do you think this is a likely scenario for a future workplace? - **[Features]** Do you think these software features are likely to be available in a future workplace?
Likelihood of actions by **Protagonist and (imagined) by Participant**: - **[Tikaani]** Do you think that someone like Tikaani would do what the Tikaani character does in this story? - **[You]** Would *you* do what the Tikaani character does in this story?
Benefit to **Company and Employee**: - **[Employee]** Do you think that this chapter's scenario would be beneficial for Tikaani? - **[Company]** Do you think that this chapter's scenario would be beneficial for the company?

Table 1: Six Likert-scale questions for each of the 13 Chapters. The same Likert scale was used for all six questions: From Very Improbable to Very Probable, plus an option "Not applicable or Decline to answer."

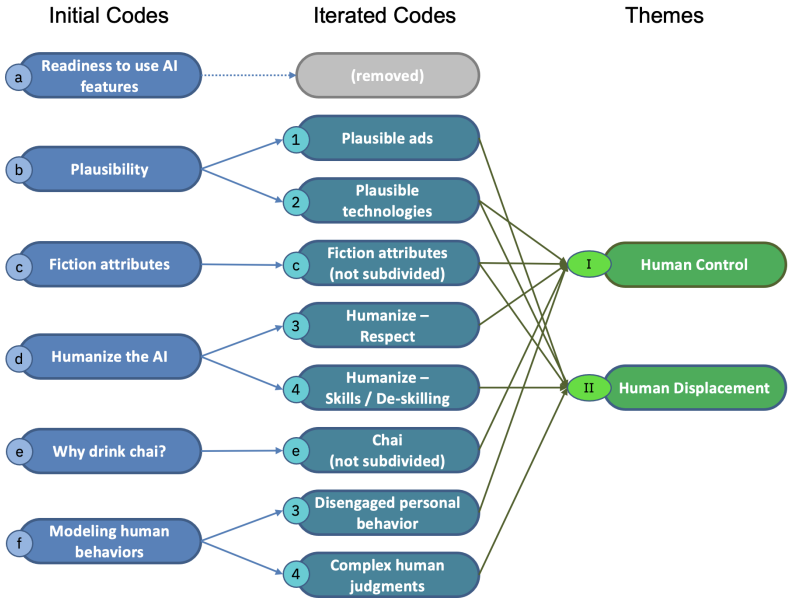


Figure 2: Initial codes, iterated codes, and themes. The codes and themes are arranged in three columns, with arrows showing connections between prior codes and successive codes or themes.

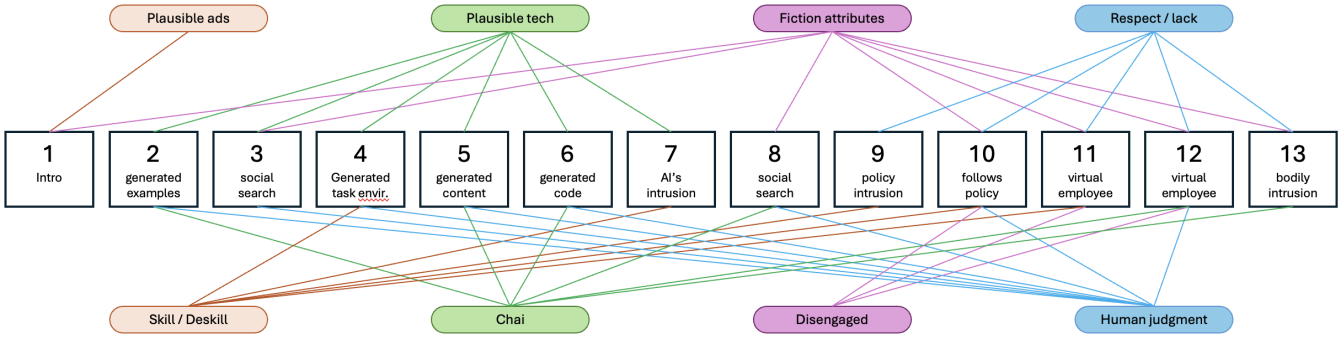


Figure 3: Eight subthemes mapped onto the 13 chapters.

For reasons of space, we focus more on the impacts of technologies and policies, and less on narrative structure. Comments on Plausible Ads and Chai, for example, occurred often, but did not contribute to the major point below.)

4.2.3 Theme: Human Control. The theme of Human Control was reflected in many participants' concerns with human agency, as opposed to algorithmic limitations on human action. P22 wanted broader options for the human to choose among: *"YPP should offer more nuanced options"* (P22, Chapter 2, Subthemes: *Respect, Skill, Human judgment*). Tensions between *digital surveillance* and *worker autonomy* surfaced in P8's concern that *"the AI... borders on micro-management."* (P8, Chapter 9, Subthemes: *Plausible tech, Fiction attributes, Respect, Skill, Disengaged, Human judgment*). Similarly, participants wrote many objections to algorithmic constraints on human imagination, such as *"diverse thinking is limited due to the limited options given"* (P26, Chapter 3, Subthemes: *Respect, Skill, Disengaged, Human judgment*) and *"[intrusive] suggestions that may even call into question the user/programmers judgement or creativity."* (P15, Chapter 4, Subthemes: *Respect, Skill, Human judgement*).

Participants were concerned about consequences of over-control by management as enforced by the AI agent, as expressed by P8: *"Seems like a losing proposition all around to me"* (P8, Chapter 13 , Subthemes: *Skill, Disengaged*). P26 provided a good summary of the tension between passive vs. active agency highlighted by Fanni et al. [34]: *"It [company policy] is trusting the AI over the human decision-maker"* (P26, Chapter 9, Subthemes: *Respect, Skill, Human judgment*). By contrast, participants argued that quality work requires human discernment and creativity. They wanted humans to govern the AI's actions, and even to be able to inactivate the AI if it interfered with humans doing their work: P6 suggested that *"Tikaani could go to manual mode and Y[PP] would be inactive"* (P6, Chapter 9, Subthemes: *Skill, Human judgment*).

4.2.4 Theme: Human Displacement. The theme of Human Displacement directly addressed a long-term value tension about humans losing their job roles, or even their employment, to a management belief that technology can do the work faster, better, cheaper, and/or under greater managerial control. P7 addressed this point directly: *"An AI assistant should not be doing developer's task."* (P7, Chapter 6, Subthemes: *Respect, Skill, Human judgment*). P23 agreed: *"Tikaani should have the authority to make the final call."* (P23, Chapter 9, Subthemes: *Respect, Skill, Human judgment*). Convergently, participants argued that an AI application would need more human assistance, *"it would take more human/AI interaction than the story implies."* (P8, Chapter 6, Subthemes: *Plausible tech, Skill, Human judgment*). If employees were *not* allowed to provide that assistance, then participants worried about deskilling: *"Tikaani's skill level is going down because they are doing less coding and I don't know how new employees could grow enough expertise."* (P26, Chapter 6, Subthemes: *Skill, Human judgment*).

Participants also highlighted the value tension in the employee-to-company relationship: *"this scenario is disrespectful to the employee"* (P15, Chapter 11, Subtheme: *Respect*) and would *"affect morale"* (P27, Chapter 11, Subthemes: *Respect, Skill, Disengaged*). Participants suggested that, with too much de-skilling, employees would prefer to work elsewhere: *"Maybe Tikaani would start looking for a different job."* (P26, Chapter 10, Subthemes: *Plausible tech, Respect, Skill, Disengaged, Human judgment*).

The continuation of P15's statement from Chapter 11 provides a good overall summary: *"This won't end well for either employee [or] company"* (P15, Chapter 11). Most participants considered that the "virtual employee" concept was a violation of employee-to-employee relationships and of company-to-employee implicit agreements.

5 DISCUSSION

As we reviewed in Section 2.1, generative AI is a technology that has potential to help people as they engage in simple and complex tasks. The strategy of *predicting the next token on the basis of preceding tokens* has proved to be a powerful aid in diverse human media projects [7, 18, 43, 71, 85, 110, 123]. Much of the commentary in the popular and academic press has been extraordinarily positive [46, 47, 92, 103]. However, like other complex technologies, generative AI vexes us - it is a "complicated, contradictory being" [130]. Applications based on generative technologies have potential to support constructive human activities, but could also be used to constrain, restrict, and even discipline human actions if they diverge from a modeled optimal next-token of human activity.

In this paper, we used a design fiction to explore both beneficial and harmful implications of generative AI for the future of work and the future of workers. Specifically, we considered the possibility that tokenization strategies could soon be applied to human work practices, in addition to the current focal attention to texts, images, music, and gestures. Through fiction, we personalized the experience of tokenization of human work through Tikaani, a diligent but perhaps increasingly disengaged tech worker. Following previous examples [60, 121, 127], we attempted to show some positive and some negative possibilities when modeling is applied to human work in general, and to the specific work practices of individual workers. Beyond expressing our own concerns, we invited tech workers - i.e., people currently working in the role of the Tikaani protagonist - to assess and comment on our concerns.

Because of their employment status, Tikanni had limited opportunities to refuse to use these implementations. In this section, we will use value tensions [91] as a way of focusing our questions and concerns around voluntary and mandated use of technologies. We will then expand on those value tensions through a discussion of refusal, non-use, and un-use. However, we believe that those concepts are in need of further refinement for situations in which usage is not voluntary, and in which decisions to use or refuse have implications beyond the individual. We therefore propose three dimensions of an initial conception of situated refusal: Who is guiding the work? How much voluntary adoption? Who else is affected by the individual worker's decisions?

5.1 Value Tensions

Theory and practice of Value Sensitive Design recommend to analyze complex technologies and policies, in part, in terms of the human and institutional stakeholders [34, 38, 100]. In the design fiction, Tikaani is an obvious stakeholder. Is YPP a second stakeholder? In immediate terms, the answer appears to be "yes." However, YPP

was made by particular human stakeholders in response to organizational priorities, and serves as a projection of those stakeholders' ideas of enablements and constraints upon Tikaani [97, 131, 139], even rendering Tikaani as an involuntary cyborg [93, 98], a "quantified employee self" [19] or a "useful body" that can be deployed "on behalf of... the corporation" [80] and its algorithmic work models, via internal virobots (chapter 13).

Feminist science and technology studies argue that all persons and organizations - all stakeholders - have a societal position, a position-based perspective (i.e., a standpoint [54, 56, 57]) and a set of values and goals based on that standpoint [52, 134]. The Drinking Chai fiction explored two distinct standpoints: that of the worker, and that of the corporation. Both Tikaani and the corporation will benefit from technologies that increase Tikaani's ability to do their job with high quality and productivity in the narratives of chapters 2, 4-7. The vignettes in these chapters represent value flows [91], with positive alignment of values and benefits to Tikaani and to the invisible management that created and deployed YPP.

By contrast, Tikaani desires greater autonomy in their work decisions, but is thwarted by YPP's enforcement of corporate policies (chapters 9-11). Tikaani also longs for collaborative contact with other workers (chapters 1, 3, 8), but YPP's preferred protocols tend toward individualistic actions and technology-mediated solutions (chapters 2, 6, 7, 11, 12). In the extreme case, Tikaani is required to become a "useful body" [80] as a partial operationalization of the corporation's control of behavior and emotional status through the virobots (chapter 13). These vignettes represent value dams [91], with poor alignment of values and benefits, and consequent losses for Tikaani and management. Participants recorded their increasing skepticism and decreasing approval of these constraining developments through their Likert ratings of the later chapters in the fiction (Section 4.2.1). Participants also critiqued these constraints through their comments, focusing on their desires for Human Control and their fears of Human Displacement by technologies (Section 4.2.3).

Value Sensitive Design can make the strongest contribution as a form of design activity [37, 38], although it also provides post-design value as an analytic tool. In a VSD analysis, we might recognize that different stakeholders experience different responsibilities and therefore make different value-based choices [34, 99, 100]. Participants considered the reduction in Human Choice to be unwelcome, and argued that Tikaani (and workers in general) should be making their own work-practice decisions.

However, Khovanskaya et al. provide a critical discussion of the organizational heritage of concepts of scientific management:

> "In Taylor's view, the difference between scientific management and 'ordinary' management was that while the latter leaves the details of work implementation up to the worker, the former places 'a great part of it upon the management.' " ([72], including a quotation from Taylor [136]).

Earlier, we mentioned the use of technology as a projection of organizational priorities upon Tikaani. This project appears to be a way to operationalize the organization's direction of the details of work implementation (see also [97]). Thus, the value dams over Human Control and even Human Displacement may reflect different stakeholders' different responsibilities - e.g., to co-workers vs. to shareholders. These felt responsibilities may reflect - or may engender - different values. Participants generally agreed with our fiction, supporting Human Choice and rejecting Human Displacement. However, in Suchman's analysis of located accountabilities [134], these values and priorities become situated in each stakeholder's social and organizational position (i.e., each person's location in the organization's hierarchy). Thus, the value dams that are highlighted in Tikaani's story, emerge from strongly held beliefs that originate in position-based identities.

When used to clarify issues in design negotiations, value flows can be used to build guarded collegiality and a calibrated form of trust, which can then strengthen a team to deal with value dams. However, in most cases of workflow technologies, the end-users have little or no presence in the decisions that informed the design of those technologies [97, 131, 139], and no ability to co-design the technologies that will affect their work, workplaces, and worklives (e.g., as contrasted with worker power in participatory design [120]). Because the hardware and software technologies are already built, there are no opportunities for Tikaani to contribute to an analysis of value flows and value dams. Tikaani can only respond to value dams through subtler actions, such as refusal, non-use, and un-use.

5.2 Refusal, Non-Use, and Un-Use

As we described earlier in Section 2.3, people have negotiated their relationships with technologies in diverse ways. Many formulations of Refusals assume that an individual is a free actor, capable of making their own choices about when to withhold data [12, 25], knowledge [24, 40, 141], and constructive engagement with the development or commerce of technologies [30, 35, 41, 67]. However, employees may face major constraints in exercising such agency [19, 72].

When refusal is considered as a kind of choice, we may fail to recognize that the ability to refuse is grounded in privilege [27, 31, 42] - i.e., "that only the privileged would be able to use it" [114]. Consumers, institutions, and some academics may be in a position to perform laudable acts of refusal, non-use, and un-use. By contrast, employees may experience radically restricted ranges of action, and often have to decide whether particular acts of refusal or resistance are worth the risk of retaliation - of being fired, or of losing opportunities to advance their careers [19, 70, 72].

Tikaani tried. When faced with the buy-out program that would have appropriated their human knowledge into technology in chapter 12, Tikaani refused [24, 40, 141] (a form of disinterest in the technology [75, 78, 115]), and Participants applauded that choice (Sections 4.2.1 and 4.2.3).

For the most part, Tikaani had to find less obvious and less complete ways of dealing with value dams. In chapter 3, YPP prompted Tikaani to make a choice based on their own individual knowledge, but Tikaani chose instead a more social-relational strategy of consulting with their peers. In chapter 8, YPP again prompted Tikaani to make an individual choice, and Tikaaani again opted to use a more structured social strategy to obtain peer input to their decision. In these ways, Tikaani adopted a social or solidarity strategy of partial or selective refusal to work in isolation ([9, 75]) despite the fact that these approaches were discouraged by the institutional rules inscribed into YPP. Tikaani was lonely in the depopulated

workplace (chapters 1, 3, 8, 11, 12), and was disenchanted with the individualistic options offered by YPP [75, 115]. In the analysis of Baumer et al., Krischkowsky et al., and Muller and Strohmayer [9, 75, 96], Tikaani was making partial use of the options that YPP offered, but was exercising their own agency about *how* to make strategized selective use of those options. Participants approved of these exercises of Human Control (Sections 4.2.1 and 4.2.3).

By contrast, Tikaani was unable to make safe use of any strategies of refusal when YPP imposed corporate rules on their use of their own knowledge in chapters 9 and 10. Tikaani had no survivable choice other than to accept the corporate discipline. Participants decried this denial of employee agency. The loss of options of refusal or non-use was even more pronounced in chapter 13, in which Tikaani was compelled to self-administer the company's surveillance technology into their own body. Participants were appalled.

5.3 Situated Refusal

In the prior section, we considered some of the limited choices available to employees, including gig workers and other workers who are not legally defined as employees, but whose livelihoods depend on asymmetric power relationships with large institutions [77]. We propose that these limitations of action may require a more nuanced perspective about refusals and non-uses, which we call *situated refusal*. Situated refusal takes a broader, more contextualized view of whether refusal, non-use, or un-use are feasible, as informed by an actor's relationship to an institution, to other actors, and to forms of collectivity.

5.3.1 Who is Guiding the Work? One dimension of situated refusal involves the type of algorithm or technology being refused, which often manifests as - or through - a policy [97]. In some cases, Tikaani was able to exercise their personal choice of action, including social and relational options that were not recommended by YPP (chapters 3 and 8). In other cases, Tikaani was required to follow policies and their instantiation as rules, which were operationalized in YPP (chapter 9). Organizations may wish to consider what can be gained through a rebalancing of control of the "details of work implementation" (e.g., [136]) with greater sharing of tactical decision-making between employee and policy-as-algorithm. For example, YPP prevented Tikaani from writing code to repair a failure in requirements (chapter 9), and Tikaani faced a bureaucratic review process and possible sanctions as a result. In a more balanced approach, the organization would voluntarily embrace certain forms of refusal, non-use, and un-use, because of the greater synergy of combining managerial knowledge and worker knowledge.

5.3.2 Degree of Voluntary Adoption. A second dimension of situated refusal depends on how control and surveillance are exercised. Tikaani's primary interaction was through YPP, which both enabled and constrained Tikaani's work. It seems obvious that YPP had to monitor Tikaani's work sufficiently to offer advice ("Recommendation follows...") and to impose restrictions ("Required action follows..."). YPP was also designed to deliver the surveillance data to management ("There is a possibility that your team-lead will need to... escalate your rationale to higher human management...") [76]. Similar issues have surfaced repeatedly among ride-service drivers [72, 77] and food workers [13], who have strategized ways to avoid managerial oversight when necessary.

However, in the closing chapter, Tikaani was required to drink the surveillance technology directly in their living body. While the virobots were a fictional construction, we note that employees are required to carry badges for location-sensing, that can both improve knowledge management [48, 145] but also track employees in workplaces [111, 122]. Mobile workers are usually required to receive assignments and report task details through mobile applications, which also provide worker data and timestamped location to organizational algorithms [72, 77]. Access to secure facilities is often mediated through required biophysical identities, such as digitized fingerprints, facial recognition, or iris pattern-matching [61], thus providing timestamped records of worker locations. Recently, employees of some companies have been required to receive implanted RFID devices in order to gain access to physical workplaces [107]. The extent of control and the extent of surveillance, as well as the extent of tranmission of those data, are in many cases determined by organizational policies rather than legal requirements. To preserve employees' autonomy and privacy (including beneficial, private actions and conversations about work [77]), organizations may wish to adopt more flexible policies about what data are recorded, who can access those data, and whether and when workers can turn off data surveillance when needed for personal or organizational reasons.

5.3.3 Collectivity and Solidarity. A third dimension of situated refusal is based in a relational context of workers who face shared challenges in a workplace. Kapoor et al. studied labor activists who were organizing online during COVID-19 [70]. They described patterns of continual re-negotiation of risk and effectiveness, sometimes organizing through institutional communication channels, but often organizing through external channels that were less subject to surveillance. Lee et al. reported similar use of non-organizational channels for organizing among Uber and Lyft drivers [77]. People in vulnerable social positions often have to strategize how to share data with other people in similar situations, without also revealing personal information to abusers or authorities [96]. The decision to refuse or to un-use previously used institutional channels was based on multiple situational factors, including collective risks, activist effectiveness, and dynamic configurations of solidarity. Thus, refusal, non-use, and un-use are not only personal decisions, but have implications and consequences for other workers and survivors. Situated refusal can involve contextual, relational, and solidarity issues.

6 CONCLUSION

In this paper, we used an extended design fiction [86, 133] to raise questions about the complexities of generative AI technologies and their implications for the future of work and the future of workers if human work becomes tokenized and modeled in the same way as data and other media. We celebrated ways in which those applications might assist and support human workers, and we critiqued ways in which those applications might harm human workers. We invited 30 colleagues to assess the fiction.

We analyzed participants' responses using quantitative and qualitative (thematic analysis [15, 137]) methods. We interpreted those

results through theories of value tensions [91], as informed by standpoint theories from Feminist science and technology studies. We considered theories of refusals, non-uses, and un-uses from the design literature [9, 10, 12, 24, 25, 40, 41, 75, 104, 141]. We noted that the ability to refuse a design is grounded in the privilege of the refuser [27, 31, 42, 114], and that future workers may not be able to engage in refusal with the same range of options as consumers or researchers. Therefore, we proposed an extension of refusal theory, with three dimensions that reflect different aspects of the complex contexts in which value tensions may occur in future workplaces.

Refusals occur in response to a value tension, and probably value dams [91]. The proposed three dimensions of situated refusal are responses to these tensions. We hope that these dimensions will help to expand the concepts of value tensions, to include tensions in authority and agency ("Who guides the work?"), the extent and methods of surveillance and control ("Degree of voluntary adoption"), and the relations among workers ("Collectivity and solidarity"). Thereby, we hope that by bringing theories of values and theories of refusal into dialog, we may strengthen both.

ACKNOWLEDGMENTS

We are grateful to Sara Berger for insightful discussions and knowledge-sharing related to refusal as privilege. We also thank Qian Pan for consultations on names and genders, and David Piorkowski, Jessica He, and Kristina Brimijoin for many discussions of conversational agents.

REFERENCES

[1] Milad Abdollahzadeh, Touba Malekzadeh, Christopher TH Teo, Keshigeyan Chandrasegaran, Guimeng Liu, and Ngai-Man Cheung. 2023. A survey on generative modeling with limited data, few shots, and zero shot. *arXiv preprint arXiv:2307.14397* (2023).

[2] Nadisha-Marie Aliman and Leon Kester. 2020. Malicious Design in AIVR, Falsehood and Cybersecurity-oriented Immersive Defenses. In *2020 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*. IEEE, 130–137.

[3] Ali Alkhatib. 2021. To live in their utopia: Why algorithmic systems create absurd outcomes. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–9.

[4] Abdulwahhab O Alshammari and Hyunggu Jung. 2017. Designing community of practice systems: A value sensitive approach. In *2017 International Conference on Informatics, Health & Technology (ICIHT)*. IEEE, 1–7.

[5] Joseph Babcock and Raghav Bali. 2021. *Generative AI with Python and TensorFlow 2: Harness the power of generative models to create images, text, and music*. Packt Publishing Ltd.

[6] Lisanne Bainbridge. 1983. Ironies of automation. In *Analysis, design and evaluation of man–machine systems*. Elsevier, 129–135.

[7] Matthew Barthet, Ahmed Khalifa, Antonios Liapis, and Georgios Yannakakis. 2022. Generative personas that behave and experience like humans. In *Proceedings of the 17th International Conference on the Foundations of Digital Games*. 1–10.

[8] Eric PS Baumer, Mark Blythe, and Theresa Jean Tanenbaum. 2020. Evaluating design fiction: The right tool for the job. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. 1901–1913.

[9] Eric PS Baumer, Jenna Burrell, Morgan G Ames, Jed R Brubaker, and Paul Dourish. 2015. On the importance and implications of studying technology non-use. *interactions* 22, 2 (2015), 52–56.

[10] Eric PS Baumer and M Six Silberman. 2011. When the implication is not to design (technology). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2271–2274.

[11] France Belanger, Mary B Burns, and Richard P Will. 1995. Transference of power: Shifting expertise to expert system developers. In *Proceedings the 11th Conference on Artificial Intelligence for Applications*. IEEE, 65–71.

[12] Ruha Benjamin. 2016. Informed refusal: Toward a justice-based bioethics. *Science, Technology, & Human Values* 41, 6 (2016), 967–990.

[13] Tiziano Bonini, Emiliano Trere, Zizheng Yu, Swati Singh, Daniele Cargnelutti, and Francisco Javier López-Ferrández. 2023. Cooperative affordances: how instant messaging apps afford. *Convergence* (2023).

[14] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.

[15] Virginia Braun and Victoria Clarke. 2022. *Common challenges in Thematic Analysis and how to avoid them*. https://youtu.be/tpWLsckpM78

[16] Edvard PG Bruun and Alban Duka. 2018. Artificial intelligence, jobs and the future of work: Racing with the machines. *Basic Income Studies* 13, 2 (2018), 20180018.

[17] Jenna Butler, Sonia Jaffe, Nancy Baym, Mary Czerwinski, Shamsi Iqbal, Kate Nowak, Sean Rintel, Abigail Sellen, Mihaela Vorvoreanu, Najeeb G Abdulhamid, Judith Amores, Reid Andersen, , Kagonya Awori, Maxamed Axmed, danah boyd, James Brand, Georg Buscher, Dean Carnignan, Martin Chan, Adam Coleman, Scott Counts, Madeleine Daepp, Adam Fourney, Daniel G Goldstein, Andy Gordon, Aaron L Halfaker, Javier Hernandez, Jake Hofman, Jenny Lay-Flurrie, Q. Vera Liao, Siân Lindley, Sathish Manivannan, Charlton Mcilwain, Sibigya Nepal, Jennifer Neville, Stephanie Nyairo, Jacki O'Niell, Victor Poznanski, Gonzalo Ramos, Nagu Rangan, Lacey Rosedale, David Rothschild, Tara Safavi, Advait Sarkar, Ava Scott, Chirag Shah, Neha Parikh Shah, Teny Shapiro, Ryland Shaw, Auste Simkute, Jina Suh, Siddharth Suri, Ioanna Tanase, Lev Tankelevitch, Adam Troy, Mengting Wan, Ryen W. White, Longqi Yang, Rrent Hecht, and Jaime Teevan. 2023. *Microsoft New Future of Work Report 2023*. https://www.microsoft.com/en-us/research/publication/microsoft-new-future-of-work-report-2023/ Accessed 7 February, 2024.

[18] Alex Calderwood, Vivian Qiu, Katy Ilonka Gero, and Lydia B. Chilton. 2020. How Novelists Use Generative Language Models: An Exploratory Study. In *Proceedings of HAI-GEN 2020 Workshop on Human-AI Co-Creation with Generative Models*.

[19] Thomas Calvard. 2019. Integrating social scientific perspectives on the quantified employee self. *Social Sciences* 8, 9 (2019), 262.

[20] Zachary Catanzaro. 2023. Beyond Incentives: Copyright in the Age of Algorithmic Production. *NYU Journal of Intellectual Property & Entertainment Law, Forthcoming* (2023).

[21] Liu Chang-de. 2006. De-skilling effects on journalists: ICTs and the labour process of Taiwanese newspaper reporters. *Canadian Journal of Communication* 31, 3 (2006), 695.

[22] Hongming Chen and Ola Engkvist. 2019. Has Drug Design Augmented by Artificial Intelligence Become a Reality? *Trends in pharmacological sciences* 40, 11 (Nov. 2019), 806–809. http://dx.doi.org/10.1016/j.tips.2019.09.004

[23] EunJeong Cheon and Norman Makoto Su. 2018. Futuristic autobiographies: Weaving participant narratives to elicit values around robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 388–397.

[24] Kimberly Christen. 2012. Does information really want to be free? Indigenous knowledge systems and the question of openness. *International Journal of Communication* 6 (2012).

[25] Marika Cifor, Patricia Garcia, T.L. Cowan, Jasmine Rault, Tonia Sutherland, "Anita Say" Chan, Jennifer Rode, "Anna Lauren" Hoffman, Niloufar Salehi, and Lisa Nakamura. 2019. *Feminist data manifest-no*. Retrieved 2023-02-09 from https://www.manifestno.com/

[26] Victoria Clarke and Virginia Braun. 2018. Using thematic analysis in counselling and psychotherapy research: A critical reflection. *Counselling and psychotherapy research* 18, 2 (2018), 107–110.

[27] Jay L Cunningham, Gabrielle Benabdallah, Daniela K Rosner, and Alex S Taylor. 2022. On the Grounds of Solutionism: Ontologies of Blackness and HCI. *ACM Transactions on Computer-Human Interaction* (2022).

[28] Alexei Czeskis, Ivayla Dermendjieva, Hussein Yapit, Alan Borning, Batya Friedman, Brian Gill, and Tadayoshi Kohno. 2010. Parenting from the pocket: Value tensions and technical directions for secure and private parent-teen mobile safety. In *Proceedings of the sixth symposium on usable privacy and security*. 1–15.

[29] P. Das, T. Sercu, K. Wadhawan, I. Padhi, S. Gehrmann, F. Cipcigan, V. Chenthamarakshan, H. Strobelt, C. dos Santos, P. Chen, Y. Yang, J. Tan, J. Hedrick, J. Crain, and A. Mojsilovic. 2020. Accelerated Antimicrobial Discovery Using Controllable Deep Generative Model and Molecular Dynamics. *Submitted to Nature, under review*. (2020).

[30] Jeffrey Dastin and Munsif Vengattil. 2020. *Microsoft bans face-recognition sales to police as Big Tech reacts to protests*. Retrieved 2023-02-09 from https://www.reuters.com/article/us-microsoft-facial-recognition-idUSKBN23I2T6

[31] Catherine D'Ignazio, Erhardt Graeff, Christina N Harrington, and Daniela K Rosner. 2020. Toward equitable participatory design: Data feminism for CSCW amidst multiple pandemics. In *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing*. 437–445.

[32] Pelle Ehn and Morten Kyng. 1987. The collective resource approach to systems design. *Computers and democracy* 17 (1987), 57.

[33] Bob Ellis and Rodney Fopp. 2001. The origins of standpoint epistemologies: Feminism, Marx and Lukacs. In *Australian Sociological Association Conference, Sydney*. Citeseer.

[34] Rosanna Fanni, Valerie Eveline Steinkogler, Giulia Zampedri, and Jo Pierson. 2020. Active Human Agency in Artificial Intelligence Mediation. In *Proceedings*

of the 6th EAI International Conference on Smart Objects and Technologies for Social Good. 84–89.
[35] J Ferguson and D Witzburg. 2021. The Chicago Police Department's Use of Shotspotter Technology. *Chicago, IL: City of Chicago Office of Inspector General* (2021).
[36] David Foster. 2019. *Generative Deep Learning*. Vol. 14. O'Reilly Media Inc., CA.
[37] Batya Friedman, Maaike Harbers, David G Hendry, Jeroen van den Hoven, Catholijn Jonker, and Nick Logler. 2021. Eight grand challenges for value sensitive design from the 2016 Lorentz workshop. *Ethics and Information Technology* 23, 1 (2021), 5–16.
[38] Batya Friedman and David G Hendry. 2019. *Value sensitive design: Shaping technology with moral imagination*. Mit Press.
[39] Anne Galloway and Catherine Caudwell. 2018. Speculative design as research method: From answers to questions and "staying with the trouble". In *Undesign*. Routledge, 85–96.
[40] Maya Indira Ganesh and Emanuel Moss. 2022. Resistance and refusal to algorithmic harms: Varieties of 'knowledge projects'. *Media International Australia* 183, 1 (2022), 90–106.
[41] Patricia Garcia, Tonia Sutherland, Marika Cifor, Anita Say Chan, Lauren Klein, Catherine D'Ignazio, and Niloufar Salehi. 2020. No: critical refusal as feminist data practice. In *Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing*. 199–202.
[42] Patricia Garcia, Tonia Sutherland, Niloufar Salehi, Marika Cifor, and Anubha Singh. 2022. No! Re-imagining Data Practices Through the Lens of Critical Refusal. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–20.
[43] Katy Ilonka Gero, Chris Kedzie, , and Lydia B. Chilton. 2020. Demo: Literary Style Transfer with Content Preservation. In *Proceedings of HAI-GEN 2020 Workshop on Human-AI Co-Creation with Generative Models*.
[44] Benyamin Ghojogh, Ali Ghodsi, Fakhri Karray, and Mark Crowley. 2021. Generative Adversarial Networks and Adversarial Autoencoders: Tutorial and Survey. *arXiv preprint arXiv:2111.13282* (2021).
[45] Jon Gillick and David Bamman. 2021. What to Play and How to Play it: Guiding Generative Music Models with Multiple Demonstrations. In *NIME 2021*. PubPub.
[46] Sharon Goldman. 2023. *How generative AI will help power your presentation in 2023*. Retrieved 2023-02-09 from https://venturebeat.com/ai/how-generative-ai-will-help-power-your-presentation-in-2023/
[47] Cindy Gordon. 2022. *Generative AI Will Continue To Accelerate In 2023: Are You Ready?* Retrieved 2023-02-09 from https://www.forbes.com/sites/cindygordon/2022/12/29/generative-ai-will-continue-to-accelerate-in-2023-are-you-ready/?sh=19c597a1650c
[48] Antonietta Grasso, Frederic Roulland, Dave Snowdon, Martin Muehlenbrock, Maurizio Mesenzani, and Roberta Morici. 2002. Supporting Informal Communication across Local and Distributed Communities. In *CSCW 2002 Workshop: Public, Community and Situated Displays*.
[49] Wayne D Gray, Bonnie E John, and Michael E Atwood. 1992. The precis of Project Ernestine or an overview of a validation of GOMS. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 307–312.
[50] Wayne D Gray, Bonnie E John, Rory Stuart, Deborah Lawrence, and Michael E Atwood. 1995. GOMS meets the phone company: Analytic modeling applied to real-world problems. In *Readings in Human–Computer Interaction*. Elsevier, 634–639.
[51] Ross Gruetzemacher, David Paradice, and Kang Bok Lee. 2020. Forecasting extreme labor displacement: A survey of AI practitioners. *Technological Forecasting and Social Change* 161 (2020), 120323.
[52] Donna Haraway. 2013. *Simians, cyborgs, and women: The reinvention of nature*. Routledge.
[53] Maaike Harbers and Mark A Neerincx. 2014. Value sensitive design of automated workload distribution support for traffic control teams. In *International Conference on Engineering Psychology and Cognitive Ergonomics*. Springer, 12–21.
[54] Sandra Harding. 2013. Rethinking standpoint epistemology: What is "strong objectivity"? In *Feminist epistemologies*. Routledge, 49–82.
[55] Sandra G Harding. 1986. *The science question in feminism*. Cornell University Press.
[56] Sandra G Harding. 2004. *The feminist standpoint theory reader: Intellectual and political controversies*. Psychology Press.
[57] Nancy Hartsock. 1983. The Feminist Standpoint: Developing the Ground for a Specifically Feminist Historical Materialism. (1983).
[58] Mark Hendrikx, Sebastiaan Meijer, Joeri Van Der Velden, and Alexandru Iosup. 2013. Procedural content generation for games: A survey. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 9, 1 (2013), 1–22.
[59] Anna Lauren Hoffmann. 2021. Even when you are a solution you are a problem: An uncomfortable reflection on feminist data ethics. *Global Perspectives* 2, 1 (2021).
[60] Stephanie Houde, Vera Liao, Jacquelyn Martino, Michael Muller, David Piorkowski, John Richards, Justin Weisz, and Yunfeng Zhang. 2020. Business (mis) Use Cases of Generative AI. *arXiv preprint arXiv:2003.07679* (2020).
[61] Anil Jain, Lin Hong, and Sharath Pankanti. 2000. Biometric identification. *Commun. ACM* 43, 2 (2000), 90–98.
[62] Mohammad Hossein Jarrahi. 2019. In the age of the smart artificial intelligence: AI's dual capacities for automating and informating work. *Business Information Review* 36, 4 (2019), 178–187.
[63] Thessa Jensen and Peter Vistisen. 2017. Ethical Design Fiction: Between storytelling and world building. *The ORBIT Journal* 1, 2 (2017), 1–14.
[64] Arif Jetha, Ali Shamaee, Silvia Bonaccio, Monique AM Gignac, Lori B Tucker, Emile Tompa, Ute Bültmann, Cameron D Norman, Cristina G Banks, and Peter M Smith. 2021. Fragmentation in the future of work: A horizon scan examining the impact of the changing nature of work on workers experiencing vulnerability. *American journal of industrial medicine* 64, 8 (2021), 649–666.
[65] Ellen Jiang, Edwin Toh, Alejandra Molina, Kristen Olson, Claire Kayacik, Aaron Donsbach, Carrie J Cai, and Michael Terry. 2022. Discovering the Syntax and Strategies of Natural Language Programming with Generative Language Models. In *CHI Conference on Human Factors in Computing Systems*. 1–19.
[66] Xisen Jin, Bill Yuchen Lin, Mohammad Rostami, and Xiang Ren. 2021. Learn continually, generalize rapidly: Lifelong knowledge accumulation for few-shot learning. *arXiv preprint arXiv:2104.08808* (2021).
[67] Elizabeth E Joh. 2017. The undue influence of surveillance technology companies in policing. *NYUL Rev. Online* 92 (2017), 19.
[68] Martin Jonsson and Jakob Tholander. 2022. Cracking the code: Co-coding with AI in creative programming education. In *Creativity and Cognition*. 5–14.
[69] Weijie Kang, Jiyang Xiao, and Junjie Xue. 2023. Generative knowledge-based transfer learning for few-shot health condition estimation. *Complex & Intelligent Systems* 9, 1 (2023), 965–979.
[70] Sayash Kapoor, Matthew Sun, Mona Wang, Klaudia Jazwinska, and Elizabeth Anne Watkins. 2022. Weaving Privacy and Power: On the Privacy Practices of Labor Organizers in the US Technology Industry. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–33.
[71] Pegah Karimi, Jeba Rezwana, Mary Lou Maher, and Nasrin Dehbozorgi. 2020. Creative Sketching Partner: An Analysis of Human-AI Co-Creativity. In *Proceedings of IUI 2020*.
[72] Vera Khovanskaya, Lynn Dombrowski, Jeffrey Rzeszotarski, and Phoebe Sengers. 2019. The Tools of Management: Adapting Historical Union Tactics to Platform-Mediated Labor. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–22.
[73] Eva Knutz and Thomas Markussen. 2014. The role of fiction in experiments within design, art & architecture-towards a new typology of design fiction. *Artifact: Journal of Design Practice* 3, 2 (2014), 8–1.
[74] Lu Kong and Jihua Li. 2022. Analysis of Production Line Project Based on Value Sensitive Design. *International Journal of Technoethics (IJT)* 13, 1 (2022), 1–11.
[75] Alina Krischkowsky, Verena Fuchsberger, and Manfred Tscheligi. 2021. Making un-use: When humans disengage with technology. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–10.
[76] Min Kyung Lee. 2018. Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society* 5, 1 (2018), 2053951718756684.
[77] Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. 2015. Working with machines: The impact of algorithmic and data-driven management on human workers. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 1603–1612.
[78] Amanda Lenhart and John B Horrigan. 2003. Re-visualizing the digital divide as a digital spectrum. *IT & society* 1, 5 (2003), 23–39.
[79] Hanlin Li, Nicholas Vincent, Janice Tsai, Jofish Kaye, and Brent Hecht. 2019. How do people change their technology use in protest? understanding. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–22.
[80] Cindy Lin and Silvia Margot Lindtner. 2021. Techniques of Use: Confronting value systems of productivity, progress, and usefulness in computing and design. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–16.
[81] Ida Lindgren. 2023. Ironies of Public Service Automation–Bainbridge Revisited. In *Proceedings of the 24th Annual International Conference on Digital Government Research*. 395–404.
[82] Yifan Liu, Zengchang Qin, Zhenbo Luo, and Hua Wang. 2017. Auto-painter: Cartoon image generation from sketch by using conditional generative adversarial networks. *arXiv preprint arXiv:1705.01908* (2017).
[83] Courtney Loder. 2014. Negotiating space between use and non-use. In *ACM CHI. A CHI 2014 Workshop Considering Why We Should Study Technology Non-use*.
[84] Ryan Louie, Andy Cohen, Cheng-Zhi Anna Huang, Michael Terry, and Carrie J Cai. 2020. Cococo: AI-Steering Tools for Music Novices Co-Creating with Generative Models.. In *HAI-GEN+ user2agent@ IUI*.
[85] Ryan Louie, Jesse Engel, and Anna Huang. 2021. Expressive Communication: A Common Framework for Evaluating Developments in Generative Models and Steering Interfaces. *arXiv preprint arXiv:2111.14951* (2021).
[86] Thomas Markussen and Eva Knutz. 2013. The poetics of design fiction. In *Proceedings of the 6th International Conference on Designing Pleasurable Products*

and Interfaces. 231–240.
[87] Wayne Martino and Kenan Omercajic. 2021. A trans pedagogy of refusal: Interrogating cisgenderism, the limits of antinormativity and trans necropolitics. *Pedagogy, Culture & Society* 29, 5 (2021), 679–694.
[88] Karl Marx. 1932. *Economic and philosophic manuscripts of 1844*. Institute of Marxism-Leninism in the Union of Soviet Socialist Republics.
[89] Joshua Meyers, Benedek Fabian, and Nathan Brown. 2021. De novo molecular design and generative models. *Drug Discovery Today* 26, 11 (2021), 2707–2715.
[90] Midjourney. 2022. Midjourney. Retrieved 07-September-2022 from https://www.midjourney.com/home/
[91] Jessica K Miller, Batya Friedman, Gavin Jancke, and Brian Gill. 2007. Value tensions in design: the value sensitive design, development, and appropriation of a corporation's groupware system. In *Proceedings of the 2007 international ACM conference on Supporting group work*. 281–290.
[92] Allison Morrow. 2023. *Welcome to the 'generative AI' era. Resistance is futile.* Retrieved 2023-02-09 from https://www.cnn.com/2023/02/08/business/nightcap-chatgpt-microsoft-google/index.html
[93] Florian Floyd Mueller, Pedro Lopes, Paul Strohmeier, Wendy Ju, Caitlyn Seim, Martin Weigel, Suranga Nanayakkara, Marianna Obrist, Zhuying Li, Joseph Delfa, et al. 2020. Next steps for human-computer integration. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–15.
[94] Michael Muller, Jeffrey Bardzell, EunJeong Cheon, Norman Makoto Su, Eric PS Baumer, Casey Fiesler, Ann Light, and Mark Blythe. 2020. Understanding the Past, Present, and Future of Design Fictions. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–8.
[95] Michael Muller and Q Vera Liao. 2017. Exploring AI Ethics and Values through Participatory Design Fictions. *Human Computer Interaction Consortium* (2017).
[96] Michael Muller and Angelika Strohmayer. 2022. Forgetting Practices in the Data Sciences. In *CHI Conference on Human Factors in Computing Systems*. 1–19.
[97] Michael Muller and Justin Weisz. 2022. Extending a Human-AI Collaboration Framework with Dynamism and Sociality. In *2022 Symposium on Human-Computer Interaction for Work*. 1–12.
[98] Kiyoshi Murata, Andrew A Adams, Yasunori Fukuta, Yohko Orito, Mario Arias-Oliva, and Jorge Pelegrin-Borondo. 2017. From a science fiction to reality: Cyborg ethics in Japan. *Acm Sigcas Computers and Society* 47, 3 (2017), 72–85.
[99] Lisa P Nathan, Batya Friedman, Predrag Klasnja, Shaun K Kane, and Jessica K Miller. 2008. Envisioning systemic effects on persons and society throughout interactive system design. In *Proceedings of the 7th ACM conference on Designing interactive systems*. 1–10.
[100] Lisa P Nathan, Predrag V Klasnja, and Batya Friedman. 2007. Value scenarios: a technique for envisioning systemic effects of new technologies. In *CHI'07 extended abstracts on Human factors in computing systems*. 2585–2590.
[101] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Z Qureshi, and Mehran Ebrahimi. 2019. Edgeconnect: Generative image inpainting with adversarial edge learning. *arXiv preprint arXiv:1901.00212* (2019).
[102] Pinchuan Ong and Ivan PL Png. 2021. Automation, Deskilling, and Labor Supply: Empirical Evidence. *Available at SSRN 3452464* (2021).
[103] Stephanie Overby. 2022. *"Artificial Intelligence (AI): 7 trends to watch for in 2022"*. "https://enterprisersproject.com/article/2022/1/artificial-intelligence-ai-7-trends-watch-2022" Accessed February 4 2022.
[104] James Pierce. 2012. Undesigning technology: considering the negation of design by design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 957–966.
[105] Ethan Plaut. 2022. Strategic illiteracies: the long game of technology refusal and disconnection. *Communication Theory* (2022).
[106] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. In *International Conference on Machine Learning*. PMLR, 8821–8831.
[107] Melanie R Rieback, Bruno Crispo, and Andrew S Tanenbaum. 2006. The evolution of RFID security. *IEEE Pervasive Computing* 5, 01 (2006), 62–69.
[108] Tapani Rinta-Kahila, Esko Penttinen, Antti Salovaara, and Wael Soliman. 2018. Consequences of discontinuing knowledge work automation-surfacing of deskilling effects and methods of recovery. (2018).
[109] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10684–10695.
[110] Steven I Ross, Fernando Martinez, Stephanie Houde, Michael Muller, and Justin D Weisz. 2023. The Programmer's Assistant: Conversational Interaction with a Large Language Model for Software Development. In *28th International Conference on Intelligent User Interfaces*.
[111] Lawrence E Rothstein. 1999. Privacy or dignity: Electronic monitoring in the workplace. *NYL Sch. J. Int'l & Comp. L.* 19 (1999), 379.
[112] Samar Sabie, Katherine W Song, Tapan Parikh, Steven Jackson, Eric Paulos, Kristina Lindstrom, Åsa Ståhl, Dina Sabie, Kristina Andersen, and Ron Wakkary. 2022. Unmaking@ CHI: Concretizing the Material and Epistemological Practices of Unmaking in HCI. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–6.
[113] Nithya Sambasivan and Rajesh Veeraraghavan. 2022. The Deskilling of Domain Expertise in AI Development. In *CHI Conference on Human Factors in Computing Systems*. 1–14.
[114] Erik Sandelin and Sarah Homewood. 2020. Design (in) actions. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. 1–9.
[115] Christine Satchell and Paul Dourish. 2009. Beyond the user: use and non-use in HCI. In *Proceedings of the 21st annual conference of the Australian computer-human interaction special interest group: Design: Open 24/7*. 9–16.
[116] Daniele Schilirò. 2023. Technological progress and the future of work. https://mpra.ub.uni-muenchen.de/117927/. (2023).
[117] Peter Senker, Joe Townsend, and Joanna Buckingham. 1989. Working with expert systems: Three case studies. *AI & society* 3 (1989), 103–116.
[118] JA Shelton. 1989. The individual 'working-to-rules': reducing determinism in Taylor-made expert systems. *Interacting with Computers* 1, 3 (1989), 338–342.
[119] Saori Shibata. 2023. 30. Deskilling and diminishing workers' autonomy in the digital workplace. *Handbook of Research on the Global Political Economy of Work* (2023), 371.
[120] Jesper Simonsen and Toni Robertson. 2013. *Routledge international handbook of participatory design*. Vol. 711. Routledge New York.
[121] Michael Skirpan and Casey Fiesler. 2018. Ad empathy: A design fiction. In *Proceedings of the 2018 ACM Conference on Supporting Groupwork*. 267–273.
[122] Leonard Sloane. 1992. Orwellian dream come true: A badge that pinpoints you. *New York Times* (1992), 14.
[123] Jason Smith. 2022. Human-AI Partnerships in Generative Music. In *International Conference on New Interfaces for Musical Expression*. PubPub.
[124] Linda Tuhiwai Smith. 2021. *Decolonizing methodologies: Research and indigenous peoples*. Bloomsbury Publishing.
[125] Jan Smits and Tijn Borghuis. 2022. Generative AI and Intellectual Property Rights. In *Law and Artificial Intelligence: Regulating AI and Applying AI in Legal Practice*. Springer, 323–344.
[126] Merlijn Smits, Chan Mi Kim, Harry van Goor, and Geke DS Ludden. 2022. From digital health to digital well-being: Systematic scoping review. *Journal of medical internet research* 24, 4 (2022), e33787.
[127] Robert Soden, Michael Skirpan, Casey Fiesler, Zahra Ashktorab, Eric PS Baumer, Mark Blythe, and Jasmine Jones. 2019. Chi4evil: Creative speculation on the negative impacts of hci research. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–8.
[128] Marie Louise Juul Søndergaard. 2020. Troubling Design: A Design Program for Designing with Women's Health. *ACM Transactions on Computer-Human Interaction (TOCHI)* 27, 4 (2020), 1–36.
[129] Marie Louise Juul Søndergaard and Lone Koefoed Hansen. 2016. PeriodShare: A bloody design fiction. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. 1–6.
[130] Marie Louise Juul Søndergaard and Lone Koefoed Hansen. 2018. Intimate futures: Staying with the trouble of digital personal assistants through design fiction. In *Proceedings of the 2018 Designing Interactive Systems Conference*. 869–880.
[131] Edward A Stohr and J Leon Zhao. 2001. Workflow automation: Overview and research issues. *Information Systems Frontiers* 3 (2001), 281–296.
[132] Taylor Stone. 2022. The streetlights are watching you: A historical perspective on value change and public lighting. *Prometheus* (2022).
[133] Miriam Sturdee and Joseph Lindley. 2019. Sketching & drawing as future inquiry in HCI. In *Proceedings of the Halfway to the Future Symposium 2019*. 1–10.
[134] Lucy Suchman. 2002. Located accountabilities in technology production. *Scandinavian journal of information systems* 14, 2 (2002), 7.
[135] Sangho Suh and Pengcheng An. 2022. Leveraging Generative Conversational AI to Develop a Creative Learning Environment for Computational Thinking. In *27th International Conference on Intelligent User Interfaces*. 73–76.
[136] Frederick Winslow Taylor et al. 1947. Scientific Management, comprising shop management. (1947).
[137] Gareth Terry, Nikki Hayfield, Victoria Clarke, and Virginia Braun. 2017. Thematic analysis. *The SAGE handbook of qualitative research in psychology* 2 (2017), 17–37.
[138] Jim Thatcher, David O'Sullivan, and Dillon Mahmoudi. 2016. Data colonialism through accumulation by dispossession: New metaphors for daily data. *Environment and Planning D: Society and Space* 34, 6 (2016), 990–1006.
[139] Andrea K Thomer and Karen M Wickett. 2020. Relational data paradigms: What do we learn by taking the materiality of databases seriously? *Big Data & Society* 7, 1 (2020), 2053951720934838.
[140] Cameron Tonkinwise. 2014. Design away. In *Design as Future-Making, Barbara Adams and Susan Yelavich (eds.). Bloomsbury Academic, London and New York*. 198–213.
[141] Eve Tuck and K Wayne Yang. 2014. Unbecoming claims: Pedagogies of refusal in qualitative research. *Qualitative Inquiry* 20, 6 (2014), 811–818.
[142] João Viera Magalhães and Nick Couldry. 2021. Giving by taking away: big tech, data colonialism and the reconfiguration of social good. *International Journal of Communication* 15 (2021), 343–362.

[143] April Yi Wang, Dakuo Wang, Jaimie Drozdal, Michael Muller, Soya Park, Justin D Weisz, Xuye Liu, Lingfei Wu, and Casey Dugan. 2021. Themisto: Towards Automated Documentation Generation in Computational Notebooks. *arXiv preprint arXiv:2102.12592* (2021).

[144] Nadine-Christine Wessel. 2023. Decision-Support Systems and Decision Making: Managing Decisional Deskilling in Human-DSS Interactions in Organizations. In *ICDS 2023 : The Seventeenth International Conference on Digital Society*.

[145] Mikael Wiberg. 2001. Knowledge management in mobile CSCW: evaluation results of a mobile physical/virtual meeting support system. In *Proceedings of the 34th Annual Hawaii International Conference on System Sciences*. IEEE, 10–pp.

[146] Chenxi Yuan and Mohsen Moghaddam. 2020. Garment design with generative adversarial networks. *arXiv preprint arXiv:2007.10947* (2020).

[147] Elaine Zibrowski, Lisa Shepherd, Kamran Sedig, Richard Booth, Candace Gibson, et al. 2018. Easier and faster is not always better: grounded theory of the impact of large-scale system transformation on the clinical work of emergency medicine nurses and physicians. *JMIR Human Factors* 5, 4 (2018), e11013.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009