

# Drinking Chai with Your (AI) Programming Partner: A Design Fiction about Generative AI for Software Engineering

Michael Muller<sup>1</sup>, Steven Ross<sup>1</sup>, Stephanie Houde<sup>1</sup>, Mayank Agarwal<sup>2</sup>, Fernando Martinez<sup>3</sup>, John Richards<sup>2</sup>, Kartik Talamadupula<sup>4</sup> and Justin D. Weisz<sup>2</sup>

<sup>1</sup>IBM Research, Cambridge MA USA

<sup>2</sup>IBM Research, Yorktown Heights NY USA

<sup>3</sup>IBM, Capital Federal, Argentina

<sup>4</sup>IBM Research, Seattle WA USA

## Abstract

Using design fiction, we develop a series of possible generative AI features and applications that could be developed in the future of humans' roles in software engineering. We use the fiction to highlight choices and value-tensions among these potential futures.

## Keywords

Generative AI; Software engineering; Workplace; Design fictions; Code-as-instance; Activity-as-instance; Employee-as-instance.

## 1. Introduction

Generative AI has the potential to improve practical work in software engineering [1, 2, 3, 4, 5, 6, 7, 8, 9]. These technologies are powerful, but there are increasing - and increasingly diverse - potential risks of applying generative AI to human work and human outcomes [10, 11, 12, 13, 14, 15]. In a paper at the 2020 HAIGEN workshop, we explored potential future societal problems with generative AI through the use of participatory design fictions [10]. In those three fictions, we invited our colleagues to speculate on possible societal harms from generative AI applications.

Here, we shift our strategy toward workplaces, and we try to take a more balanced view, considering both potential benefits and potential risks of generative applications in workplaces of the future. Workplace adaptations due to the COVID pandemic have accelerated sociotechnical trends of changed work-practices and changed technolog-

ical infrastructures. We adopt a different strategy based on design fictions as a research method [16, 17] that allows us to consider a set of related benefits and risks of generative AI applications that might be used in future workplaces.

We contribute

- an interrelated set of speculations regarding future generative AI applications in workplaces, and
- a consideration of some value tensions that may emerge between employee needs and organizational needs

We also critique our work as being *half-done*, and we describe possible ways to complete the work in the near future.

## 2. Background

### 2.1. Generative AI for Software Engineering

One of the strengths of generative algorithms and applications is their ability to create instances from a "learned" class of examples, including projects that involve images [18, 19, 20], videos [21], music [22, 23], molecules [24, 25], texts of many types [26, 27, 28, 29, 30], and diverse other media and categories (e.g., [31, 32, 33]). When analyzed as sequences of tokens, these "learned" patterns can function as predictions of (e.g.) the next word in a text or a software program; this "next-in-

Joint Proceedings of the ACM IUI Workshops 2022, March 2022, Helsinki, Finland

✉ michael\_muller@us.ibm.com (M. Muller); steven\_ross@us.ibm.com (S. Ross); stephanie.houde@ibm.com (S. Houde); Mayank.Agarwal@ibm.com (M. Agarwal); martferc@ar.ibm.com (F. Martinez); ajtr@us.ibm.com (J. Richards); krtalamad@us.ibm.com (K. Talamadupula); jweisz@us.ibm.com (J. D. Weisz)

🆔 0000-0001-7860-163X (M. Muller); 0000-0002-2533-9946 (S. Ross); 0000-0002-0246-2183 (S. Houde); 0000-0002-8442-2651 (M. Agarwal); 0000-0001-7172-4805 (F. Martinez); 0000-0001-8489-2170 (J. Richards); 0000-0002-4628-3785 (K. Talamadupula); 0000-0003-2228-2398 (J. D. Weisz)

© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).  
CEUR Workshop Proceedings (CEUR-WS.org)

sequence” aspect is a major focus of this paper, here extended to the “next” user actions.

While there is not space in this position paper to review the rich history of generative algorithms, we note briefly that many reports describe generative AI applications with both potential benefits [34, 35, 36] and potential risks [37, 38, 39, 14, 12, 10] to individuals or societies. Our team have been studying how generative application scan provide significant value to software engineers [26, 8, 28, 40], and our IUI 2022 paper reports measurable successes [27]. In this position paper, we ask what types of instances might be created in the future, and whether organizations or societies may wish to place limits on the generation of certain types of instances.

In the history of theorizing *mixed-initiative* dynamics between humans and AIs, most scholars have imagined batch-like processes in which one party (human or AI) takes a first major step, and the second party (AI or human) completes the work through a second major step [41, 42, 43, 44, 45, 46, 47, 48]. While Conversational User Interfaces (CUIs) provide limited forms of on-going interactions, the sheer computational requirements of generative AI algorithms have made it difficult to design and build generative dialogic applications until recently (e.g., [49, 50, 51]). An important problem in these dialogs will be how to control the generative AI execution from moment-to-moment - i.e., how to “tune” [52, 53] or “steer” [22, 23] the algorithm toward outcomes of value to their human partners. We speculate on how possible control-parameters might be provided to human users, and how particular values of those parameters might be recommended through content-based or social-based algorithms.

## 2.2. Design Fictions

For more than a decade, scholars have used design fictions as a core research method [16, 17]. Based on theorizing by Peirce [54] over a century ago, Dunne and Martin summarized the potential of design fictions as *abductive* methods, stating that

“The designers who can solve the most wicked problems do it through collaborative integrative thinking, using *abductive* logic, which means the logic of *what might be*. Conversely, *deductive* and *inductive* logic are the logic of what *should be* or *what is...* [55] (italics added)

Abductive reasoning has also been claimed as a major tool for creatively building theory in approaches such as grounded theory [56, 57, 58, 59] and thematic analysis [60, 61]. We apply future-oriented design fictions to extrapolate current phenomena and trends into possible futures [62, 63] and to begin to interrogate those possible

futures for their human implications and impacts [64, 65]. We hope that design fictions can contribute to the mixed-initiative discussions mentioned earlier, focusing attention the types of relationships that we envision between humans and AIs [47, 66].

Design fictions may take several forms, including text, images, film, video, theatre, and physical objects [67, 68, 69, 70, 71, 17], and may be particularized into recognizable genres such as fictional job adverts [72], enactments [73], technology probes [74, 75, 76], product catalogs [74], autobiographies [77], and even clearly-fictitious conference papers [78, 79]. More familiarly, design fictions can take the form of stories [10, 66, 80, 81, 82, 83, 84, 85, 86, 71, 87, 88, 89]. We adopt the method of telling a story, through which we explore trends and possibilities at the intersection of technology, skilled human work, and values.<sup>1</sup>

## 3. Design Fiction

### 3.1. Crafting the Design Fiction

Design fiction scholars emphasize that a fiction should provide a *perceptual bridge* [90] between the reader and novel concepts that may be challenging [64, 91] or even upsetting [74, 66, 10]. The temporary suspension of disbelief is considered important [92, 90], and this is accomplished through narrative integrity and consistency of the imagined world [85, 93] and through empathy with its protagonists [94, 95]. Auger states that “careful management of the speculation” is important, because “if it strays too far... the audience will not relate to the proposal resulting in a lack of engagement or connection” [90]. Design fictions use methods such as *diegesis* (the creation of a story-world) to communicate new possibilities and to discuss their consequences [96, 97, 16].

We apply these principles by setting our story in a recognizable software engineering workplace - albeit with futuristic tools and the gradual unveiling of an unsettling workplace culture. To avoid aspects of the “uncanny valley” that can occur when AI agents behave in eerily human ways [98, 15], we designed the conversational style of the generative AI assistant - “Your Programming Partner (YPP)” - to be recognizably non-human - even robotic in some ways.

### 3.2. Strategic Ambiguity

While design fictions are usually crafted for specific intentions, it is also important that they act as *cul-*

<sup>1</sup>Methodological treatments of how to approach or write design fictions may be found in Markussen and Knutz [16], Sturdee et al. [17], Huusko et al. huusko2018structuring, Blythe and Wright [65], and Cheon and Su [77]; see also Blythe and Encinas [68] and Baumer et al. [78] for approaches to assessing or evaluating design fictions.

*tural probes* [99], actively engaging the reader to form their own interpretations and to draw their own conclusions [100, 101, 75, 16, 102]. Coulton et al. note that "it seems that Design Fiction has ambiguity 'baked in'" [103], and researchers have tried to balance among diegetic factors such as familiarity and consistency, vs. the ambiguity that is needed to encourage new (e.g., abductive [55]) thinking among their readers [66, 68, 104].

The strategy of ambiguity can be particularly important if the fiction addresses value tensions [105, 84, 80]. Ambe [83] and Huusko [64] argue that many technology applications may have both utopian and dystopian implications, and these implications may be different for different stakeholders [106, 107]. Feminist technoscience convergently urges us to consider each person's perspective, and to question power differentials [108, 109, 110, 111]. We noted above that design fictions may be designed to raise challenging or upsetting questions. Using a formulation from Haraway's feminist theorizing [112], Sondergaard et al. suggest that we need to "stay with the trouble" of value tensions in AI:

"Might we allow the [AI assistant] 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." [70]

We tried to apply these rather open-ended concepts by deliberately leaving certain details unspecified, including the gender-identities of the actors, the nature of the actors in certain workplace roles, and the reasons for the loneliness that is experienced by the protagonist. We also presented certain potential value tensions in a relatively neutral way, to emphasize the questions rather than to impose premature closure on complex ethical topics.

For brevity, we now proceed directly to the design fiction.

Fiction	Notes
<b>Drinking Chai with Your (AI) Programming Partner</b>	
<p>&lt;1&gt; 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 it 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).</p>	<p>Motivation and Setting.</p> <p>Problems with legacy applications [26, 27].</p> <p>Why is the office so empty?</p> <p>Introduce the AI Assistant.</p>
<p>&lt;2&gt; 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..."</p>	<p>Establish YPP's robotic voice, reducing the likelihood of uncanny valley effects [113, 114]. Messages from multiple entities may be delivered through the YPP conversational interface.</p>
<p>&lt;3&gt; 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.</p>	<p>Personalization of the social presence of the AI.</p> <p>Speech-to-text.</p> <p>Gestural component of UI.</p>
<p>&lt;4&gt; "Drink chai," said YPP. Tikaani took a sip of chai.</p>	<p>Introduce a theme that will become complexified later.</p>
<p>&lt;5&gt; 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.</p>	<p>Combined nature of recommendations as both informational and actionable.</p>

---

**Fiction**

---

---

**Notes**

---

<6> 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.

<7> YPP displayed, “Further guidance is required. Recommendations follow. (1) Granularity of microservices can be one of: 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.

<8> YPP asked, “You usually prefer a JupyterGen notebook rather than VSEdit. Decision required: Should a JupyterGen notebook with your usual defaults be used for this project?”

<9> 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.

Human can override AI’s recommendation...

... and can specify a different, more social way of choosing data.

Introduce Yu as a colleague, to be complexified below.

Similarly to <5>, each row of the display functions as both informational and actionable.

Build YPP’s robotic nature through engineering-style passive voice. Show multiple “tunable parameters” of the generative algorithm [115, 116, 117].

Human can revise recommendation or accept as-is.

YPP has “learned” a generative model of Tikaani’s actions from past projects, and has predicted Tikaani’s most likely “next action” today, saving Tikaani’s time in the notebook.

Generative AI provides partial results for human to complete [26].

Fiction	Notes
<p>&lt;10&gt; 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.</p>	<p>Reflection on how the human's role has changed - mostly for the better.</p> <p>AI flags low-confidence outcomes for human action.</p> <p>User modeling.</p> <p>Benefits of human+genAI collaboration.</p>
<p>&lt;11&gt; 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, "CyberHR hope that this recognition will help you to have a good day, a very good day."</p>	<p>Incentives delivered through the conversational UI.</p>
<p>&lt;12&gt; "Drink chai," YPP said. Tikaani took a sip of chai.</p>	
<p>&lt;13&gt; 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 <code>QuikClientPortfolio()</code> from commercial package <code>QuikFinance</code> would be 20% faster, but has license fees that are usually \$5k to \$8k. Do you want to use <code>QuikClientPortfolio()</code> and submit a request for approval?"</p>	<p>AI initiates consultation, but allows human to make the business decision.</p>
<p>&lt;14&gt; Unsure of how to make the decision, Tikaani typed, "Inquire through YourPartnetNet for colleagues who have knowledge of <code>QuikClientPortfolio</code>. Poll 'Which release of <code>QuikClientPortfolio</code> is stable?' and indicate my status as Blocked. Use probable-knowldege feature. Anonymous responses are acceptable."</p>	<p>The network of "Partner" AIs functions as a Transactional Memory System (TMS) [118]. It stores records of "who knows what" [119, 120], and it "learns" each employee's knowledge trajectory, so that it can generate probable knowledge-states based on past activities [121].</p>
<p>&lt;15&gt; 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."</p>	
<p>&lt;16&gt; Tikaani decided to accept the majority opinion. They typed, "Estimate likelihood of approval for <code>QuikClientPortfolio</code> Release 15.3, based on data from the current year only."</p>	
<p>&lt;17&gt; YPP displayed, "Estimation follows. Package <code>QuikFinance</code> 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 <code>QuikClientPortfolio()</code> and submit a request for approval?"</p>	<p>AI can access the human's personal history of this category of request.</p> <p>Reinforce the robotic-voice repetition of the question.</p>

Fiction	Notes
<p>&lt;18&gt; Tikaani pressed the speech-to-text key and said, "Paste Quik-ClientPortfolio, 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.</p>	<p>Generative production of request-for-approval.</p> <p>Generative AI provides partial results for the human to complete [26].</p>
<p>&lt;19&gt; "Drink chai," YPP said. Tikaani sipped.</p>	
<p>&lt;20&gt; 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 there would need to be additional factors added to the model. They began to rewrite the generated 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.</p>	<p>Human contextual knowledge is uniquely informative [122, 123].</p> <p>However, the human is violating work norms.</p> <p>The human is stubborn.</p>
<p>&lt;21&gt; 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."</p>	<p>AI explains the costs of the human action...</p> <p>... and the approvals process that the human must follow...</p> <p>... including a policy-based risk.</p>
<p>&lt;22&gt; 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.</p>	<p>AI has become a projection of executive policy and power [124].</p>
<p>&lt;23&gt; 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.</p>	<p>Did the CyberHR AI use emotion-sensing algorithms [125] to determine that the human may be persuadable to take the buy-out that will benefit the company? [126].</p> <p>The same generative technology that allowed YPP to predict the human's preferred JupyterGen implementation environment &lt;8&gt;, can be used to create v-Tikaani.</p>

Fiction	Notes
<p>&lt;24&gt; After the frustration of having to justify a necessary rewrite, Tikaani was almostly 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. Tikaani declined the offer, and went back to coding.</p>	<p>Yu had agreed to be replaced by a virtual employee that was generatively modeled based on Yu’s work records.</p>
<p>&lt;25&gt; “Drink chai,” YPP said. Tikaani glared at the screen.</p>	<p>Tikaani is experiencing negative affect.</p>
<p>&lt;26&gt; “Drink chai, YPP repeated, “You are at 46%. Explanation follows. Having a good day is an employee responsibility. Chai helps. Below 40%, level-1 reporting to management is mandated. Required action follows. Drink chai.” Was that a robot joke? Tikaani had heard rumors of an experimental emotion module. Just before Yu had retired and been replaced with v-Yu, Yu had suggested that –</p>	<p>Drinking chai is required by policy, and is sensed by the AI.</p>
<p>&lt;27&gt; YPP interrupted Tikaani’s reverie. “Drink chai. You are at 42%.”</p>	
<p>&lt;28&gt; Tikaani gulped down a full mouthful of chai, and braced against its 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 the short-lived CHAI virobots were linked to the cyberHR department through near-field communications. The virobots monitored 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 their blood-brain barrier to reach the emotion centers of the limbic system. YPP’s warnings had left Tikaani is 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.</p>	<p>Similarly to policy-based campaigns to improve employee engagement[127, 128, 129], the company helps employees to work happily and productively through biological interventions...</p> <p>... which are highly effective, perhaps benefiting company productivity and perhaps employee mental health?</p>

Table 1: Design Fiction with explanatory notes.



## 4. Discussion

For brevity, we will link our Discussion points to the fiction by reference to numbered paragraphs - e.g., "<3>."

### 4.1. Generative AI Applications and Features in Software Engineering Workplaces

The "Drinking Chai" story explored a series of increasingly futuristic applications of generative AI. We briefly review them here:

- We began with the well-understood domain of **generative software translation** in paragraphs <1, 3, 5, 7, 8, 9, 10>, and we emphasized the need for human-AI partnerships in translation <10, 13, 16, 20> [28, 26, 27, 130], sometimes guided by the AI's flagging of low-confidence translations <10> [26, 27] and by the AI's recommendation of alternate classes and libraries <10, 13>. Applied inflexibly, the principle of generative anticipation of the "best" coding usages will lead to trouble in paragraphs <20-22>.
- We also included currently-available capabilities to **generate documentation** for the generated and modified code, as has been done in limited ways for source code <18> [40, 8] and also for certain sub-genres of journalism (e.g., [131]).
- We proposed capabilities for the human to **control, steer, and tune** specific aspects of the generative processes and outcomes <5, 7> [115, 22, 132].
- We included GUI **style transfer** as a way of refining the generic translations <6> [133].
- We proposed a more futuristic capability to **parse a specification document** into an architecture plan, and then into a high-level class structure <3> with necessary human guidance <5, 7> .
- Further into the future, we proposed that a generative assistant could **learn the work-practices of its human partner**, and could save the human's effort by suggesting and then implementing the anticipated "next steps" <8, 9, 17>.
- More controversially, we considered that organizations could develop **virtual versions of particular employees** through "learning" their individual work-practices and then implementing those patterns into a virtual replacement for the employee <23, 24, 26>. We acknowledge that this idea is futuristic. GANs and related algorithms require large amounts of data. While organizations may be able to use existing generative algorithms for the general case of human actions, further

research will be required to specialize these patterns to an individual human's pattern through a smaller set of personal data. We leave open the question of whether this would be a desirable outcome.

- Again controversially, we considered that organizations might use generative technologies to "learn" a "**best**" set of **employee attributes**, and might seek to impose those attributes on less-compliant employees <28>.

In the preceding discussion, we listed an escalating series of generative features, and we provided evidence (where it exists) that current research may be trending toward those features. It may be useful to address the plausibility of the concept of virobots to influence employees' emotions <26, 28>. While the control aspects are futuristic, we note that in-dwelling digital devices have been part human medical audiology, endocrinology, and neurology for at least a decade [134], now as part of a medical approach to Internet of Things [135] with concomitant privacy and security risks [136, 137]. Some of these devices are already being used to modify the patient's brain state [138, 139]. In the commercial space, Applied Digital company's Verichip is marketed as a subdermal injection of an RFID tag for building access[140].

Further, we note that organizations have for years played white noise [141] or curated music [142, 143, 144, 141] to affect employees' mental status for specific organizational purposes such as enhanced concentration [143], job-engagement [142], and even reduction in costly *employee-controlled* overtime hours [144]. Thus, employers' interests in affecting employee mental state are already part of conventional office technologies, and implantable digital devices are already a medical and industrial reality, as is the use of signals from these devices to modify a person's internal milieu. The only questions that separate the existing state-of-the-art technologies and our paragraphs <26, 28> are: *What kinds of psychological modifications might be implemented in the future?* and *How would they be "delivered?"* and *Who controls those modifications?* and of course *Is this a desirable future?*

### 4.2. Value Tensions

In keeping with theory and practice of design fictions [83, 64, 107, 91] and feminist technoscience [109, 110, 108, 111], we hoped to raise questions of personal, organizational, and societal values through our fiction. Our use of the principle of strategic ambiguity [103, 106, 70] helps to highlight some of the tensions:

- What are the trade-offs of panoptic [145, 146] surveillance on employees' work-practices to achieve organizational goals of productivity <8,

- 10, 13, 17, 20-22>? Can we distinguish between generative assistance and generative intrusion?
- When should work be governed by choice vs. organizationally-determined "best" practices <20-22>? How can we balance between "norms" and individual and group innovations <28>?
- If AI agents serve as "teammates" [47], then what are the appropriate human-AI collaboration dynamics [42, 41, 43, 44, 45]? Should they assist <1, 3, 6-10>, advise <13, 17>, monitor <21>, and/or sanction employees <21, 26-28>? How do we adjudicate competing claims about replacing human employees by digital employees <23-24>?
- Who should own the intellectual property rights of the data used to model the digital employee <23>?
- Is Tikaani's team-lead human or algorithmic <21>? Managers often interpret, modify, and implement organizational policies to their employees. How would that managerial function be different if an algorithm were to operationalize those policies without human considerations?
- Why is Tikaani so concerned with loneliness <1, 6, 28>?

We hope that these tensions will help us to "stay with the trouble" [70, 112, 108] about how generative AI might affect employees and organizations, and under whose guidance.

## 5. Conclusion

We have applied design fictions as a research method [16, 17] to develop conjectures about possible futures of generative AI features and applications, and to raise values-based questions about those possible futures. It is fair to ask, "what did we learn from this research method?" In the Discussion, we explored two topics.

We considered current generative AI approaches, and we projected them into possible futures. Generative software translation and generative documentation are current capabilities [28, 26, 27, 130, 40, 8, 131]. Further, we note that there is already research under way to provide finer controls over generative algorithms citelouie2020novice, louie2020cococo, zhou2020generative - although there is much work yet to be done. We used those concepts as "starting points," to explore more futuristic ideas.

If multiple generative AI applications predict the "next token" in a sequence of tokens, then we speculated about the nature of possible future tokens, and what those tokens might be used for. We considered what might be possible - for good or ill - if human actions were treated as be tokens. We were then able to imagine helpful scenarios, in which an AI could beneficially anticipate the

next human action, and could prepare for that action. We also imagined what might be harmful scenarios, in which a human actor might be replaced by a token-based model trained on that person's history of actions.

Finally, based in feminist technoscience [109, 110, 108, 111] and value sensitive design [147], we explored implicit values, and the likelihood of value tensions among diverse stakeholders in each of these possible futures.

Despite the design fiction tradition of publishing a fiction without empirical data [65, 81, 86, 71, 126, 87, 66], we want to broaden the conversation. Our next steps will involve participatory design fiction methods' [148, 149, 150, 89, 80] to make the work more polyvocal [101] and more reflective of our technical community's diverse opinions and aspirations.

## References

- [1] A. Hindle, E. T. Barr, Z. Su, M. Gabel, P. Devanbu, On the naturalness of software, in: 2012 34th International Conference on Software Engineering (ICSE), IEEE, 2012, pp. 837–847.
- [2] V. Raychev, M. Vechev, A. Krause, Predicting program properties from "big code", in: Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, 2015, pp. 111–124.
- [3] M. Bruch, M. Monperrus, M. Mezini, Learning from examples to improve code completion systems, in: Proceedings of the 7th joint meeting of the European software engineering conference and the ACM SIGSOFT symposium on the foundations of software engineering, 2009, pp. 213–222.
- [4] A. Svyatkovskiy, S. K. Deng, S. Fu, N. Sundaresan, Intellicode compose: Code generation using transformer, arXiv preprint arXiv:2005.08025 (2020).
- [5] M. Tufano, D. Drain, A. Svyatkovskiy, S. K. Deng, N. Sundaresan, Unit test case generation with transformers, arXiv preprint arXiv:2009.05617 (2020).
- [6] X. Guo, Towards automated software testing with generative adversarial networks, in: 2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks - Supplemental Volume (DSN-S), IEEE Computer Society, Los Alamitos, CA, USA, 2021, pp. 21–22. URL: <https://doi.ieeecomputersociety.org/10.1109/DSN-S52858.2021.00021>. doi:10.1109/DSN-S52858.2021.00021.
- [7] L. Moreno, J. Aponte, G. Sridhara, A. Marcus, L. Pollock, K. Vijay-Shanker, Automatic generation of natural language summaries for java classes, in: 2013 21st International Conference on

- Program Comprehension (ICPC), IEEE, 2013, pp. 23–32.
- [8] A. Y. Wang, D. Wang, J. Drozdal, M. Muller, S. Park, J. D. Weisz, X. Liu, L. Wu, C. Dugan, Themisto: Towards automated documentation generation in computational notebooks, arXiv preprint arXiv:2102.12592 (2021).
- [9] B. Roziere, M.-A. Lachaux, L. Chausson, G. Lample, Unsupervised translation of programming languages., in: NeurIPS, 2020.
- [10] S. Houde, V. Liao, J. Martino, M. Muller, D. Piorkowski, J. Richards, J. Weisz, Y. Zhang, Business (mis) use cases of generative ai, arXiv preprint arXiv:2003.07679 (2020).
- [11] L. A. Liikkanen, It ain't nuttin' new—interaction design practice after the ai hype, in: IFIP Conference on Human-Computer Interaction, Springer, 2019, pp. 600–604.
- [12] N.-M. Aliman, L. Kester, Malicious design in aivr, falsehood and cybersecurity-oriented immersive defenses, in: 2020 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR), IEEE, 2020, pp. 130–137.
- [13] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, et al., On the opportunities and risks of foundation models, arXiv preprint arXiv:2108.07258 (2021).
- [14] K. McGuffie, A. Newhouse, The radicalization risks of gpt-3 and advanced neural language models, arXiv preprint arXiv:2009.06807 (2020).
- [15] L. Whittaker, T. C. Kietzmann, J. Kietzmann, A. Dabirian, “all around me are synthetic faces”: The mad world of ai-generated media, IT Professional 22 (2020) 90–99.
- [16] T. Markussen, E. Knutz, The poetics of design fiction, in: Proceedings of the 6th International Conference on Designing Pleasurable Products and Interfaces, 2013, pp. 231–240.
- [17] M. Sturdee, J. Lindley, Sketching & drawing as future inquiry in hci, in: Proceedings of the Halfway to the Future Symposium 2019, 2019, pp. 1–10.
- [18] Y. Liu, Z. Qin, Z. Luo, H. Wang, Auto-painter: Cartoon image generation from sketch by using conditional generative adversarial networks, arXiv preprint arXiv:1705.01908 (2017).
- [19] K. Nazeri, E. Ng, T. Joseph, F. Z. Qureshi, M. Ebrahimi, Edgeconnect: Generative image inpainting with adversarial edge learning, arXiv preprint arXiv:1901.00212 (2019).
- [20] P. Karimi, J. Rezwana, M. L. Maher, N. Dehbozorgi, Creative sketching partner: An analysis of human-ai co-creativity, in: Proceedings of IUI 2020, 2020.
- [21] J. Babcock, R. Bali, Generative AI with Python and TensorFlow 2: Harness the power of generative models to create images, text, and music, Packt Publishing Ltd, 2021.
- [22] R. Louie, A. Cohen, C.-Z. A. Huang, M. Terry, C. J. Cai, Cococo: Ai-steering tools for music novices co-creating with generative models., in: HAI-GEN+ user2agent@ IUI, 2020.
- [23] R. Louie, J. Engel, A. Huang, Expressive communication: A common framework for evaluating developments in generative models and steering interfaces, arXiv preprint arXiv:2111.14951 (2021).
- [24] 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, A. Mojsilovic, Accelerated antimicrobial discovery using controllable deep generative model and molecular dynamics, Submitted to Nature, under review. (2020).
- [25] H. Chen, O. Engkvist, Has Drug Design Augmented by Artificial Intelligence Become a Reality?, Trends in pharmacological sciences 40 (2019) 806–809. URL: <http://dx.doi.org/10.1016/j.tips.2019.09.004>.
- [26] J. D. Weisz, M. Muller, S. Houde, J. Richards, S. I. Ross, F. Martinez, M. Agarwal, K. Talamadupula, Perfection not required? human-ai partnerships in code translation, in: 26th International Conference on Intelligent User Interfaces, 2021, pp. 402–412.
- [27] J. D. Weisz, M. Muller, S. I. Ross, F. Martinez, S. Houde, M. Agarwal, K. Talamadupula, J. Richard, Better together? an evaluation of ai-supported code translation, IUI 2022 Workshop Human AI with Generative AI (2022).
- [28] K. Talamadupula, Applied ai matters ai4code: Applying artificial intelligence to source code (2021).
- [29] A. Calderwood, V. Qiu, K. I. Gero, L. B. Chilton, How novelists use generative language models: An exploratory study, in: Proceedings of HAI-GEN 2020 Workshop on Human-AI Co-Creation with Generative Models, 2020.
- [30] K. I. Gero, C. Kedzie, L. B. Chilton, Demo: Literary style transfer with content preservation, in: Proceedings of HAI-GEN 2020 Workshop on Human-AI Co-Creation with Generative Models, 2020.
- [31] A. Harsuvanakit, Elbo chair, 2016. URL: <https://gallery.autodesk.com/fusion360/projects/elbo-chair--generated-in-project-dreamcatcher-made-with-fusion-360>.
- [32] B. Quanz, W. Sun, A. Deshpande, D. Shah, J.-e. Park, Machine learning based co-creative design framework, arXiv preprint arXiv:2001.08791 (2020).
- [33] Autodesk, Reimagining the future of air travel, 2016. URL: <https://www.autodesk.com/customer-stories/airbus>.

- [34] S. Ali, D. Parikh, Telling creative stories using generative visual aids, arXiv preprint arXiv:2110.14810 (2021).
- [35] M. Mirka, M. France-Pillois, G. Sassatelli, A. Gamatié, A generative ai for heterogeneous network-on-chip design space pruning, in: 25th Design, Automation, and Test in Europe Conference (DATE 2022), 2022.
- [36] J. Meyers, B. Fabian, N. Brown, De novo molecular design and generative models, *Drug Discovery Today* 26 (2021) 2707–2715.
- [37] E. M. Bender, T. Gebru, A. McMillan-Major, S. Shmitchell, On the dangers of stochastic parrots: Can language models be too big?, in: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 2021, pp. 610–623.
- [38] Y. Mirsky, A. Demontis, J. Kotak, R. Shankar, D. Gelei, L. Yang, X. Zhang, W. Lee, Y. Elovici, B. Biggio, The threat of offensive ai to organizations, arXiv preprint arXiv:2106.15764 (2021).
- [39] M. K. Land, J. D. Aronson, Human rights and technology: New challenges for justice and accountability, *Annual Review of Law and Social Science* 16 (2020) 223–240.
- [40] M. Muller, A. Y. Wang, S. Ross, J. D. Weisz, M. Agarwal, K. Talamadupula, S. Houde, F. Martinez, J. Richards, J. Drozdal, X. Liu, D. Piorkowski, D. Wang, How data scientists improve generated code documentation in jupyter notebooks, 2021. URL: <https://hai-gen2021.github.io/program/>.
- [41] J. A. Biles, Genjam: Evolution of a jazz improviser, in: *Creative evolutionary systems*, Elsevier, 2002, pp. 165–187.
- [42] P. M. Fitts, Human engineering for an effective air-navigation and traffic-control system. (1951).
- [43] T. B. Sheridan, M. K. Tulga, A model for dynamic allocation of human attention among multiple tasks, in: *Proceedings of the 14th Annual Conference on Manual Control*, 1978, pp. 569–592.
- [44] R. Parasuraman, Designing automation for human use: empirical studies and quantitative models, *Ergonomics* 43 (2000) 931–951.
- [45] B. Shneiderman, Human-centered artificial intelligence: Trusted, reliable & safe, arXiv preprint arXiv:2002.04087 (2020).
- [46] E. Horvitz, Principles of mixed-initiative user interfaces, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '99*, Association for Computing Machinery, New York, NY, USA, 1999, p. 159–166. URL: <https://doi.org/10.1145/302979.303030>. doi:10.1145/302979.303030.
- [47] I. Seeber, E. Bittner, R. O. Briggs, T. De Vreede, G.-J. De Vreede, A. Elkins, R. Maier, A. B. Merz, S. Oestereiß, N. Randrup, et al., Machines as teammates: A research agenda on ai in team collaboration, *Information & management* 57 (2020) 103174.
- [48] T. O'Neill, N. McNeese, A. Barron, B. Schelble, Human–autonomy teaming: A review and analysis of the empirical literature, *Human Factors* (2020) 0018720820960865.
- [49] P. Pataranutaporn, V. Danry, J. Leong, P. Pungpongson, D. Novy, P. Maes, M. Sra, Ai-generated characters for supporting personalized learning and well-being, *Nature Machine Intelligence* (2021) 1–10.
- [50] O. Nilay Yalcin, N. Abukhodair, S. DiPaola, Empathic ai painter: A computational creativity system with embodied conversational interaction, arXiv e-prints (2020) arXiv–2005.
- [51] Ö. N. Yalçın, N. Abukhodair, S. DiPaola, Empathic ai painter: A computational creativity system with embodied conversational interaction, in: *NeurIPS 2019 Competition and Demonstration Track*, PMLR, 2020, pp. 131–141.
- [52] M. van Doorn, S. Duivestijn, D. Mamtani, T. Pepping, Infinite machine creativity (n.d.). URL: <https://labs.sogeti.com/wp-content/uploads/2020/06/SogetiLabs-Infinite-Machine-Creativity.pdf>, accessed 1 January 2022.
- [53] D. Friedmana, D. Pollaka, Image co-creation by non-programmers and generative adversarial networks (2021). URL: <http://ceur-ws.org/Vol-2903/IUI21WS-HAIGEN-4.pdf>, accessed 1 January 2022.
- [54] C. S. Peirce, The essential Peirce: Selected philosophical writings, vol. 2, Indiana University Press, 1903.
- [55] D. Dunne, R. Martin, Design thinking and how it will change management education: An interview and discussion, *Academy of Management Learning & Education* 5 (2006) 512–523.
- [56] S. Timmermans, I. Tavory, Theory construction in qualitative research: From grounded theory to abductive analysis, *Sociological theory* 30 (2012) 167–186.
- [57] I. Tavory, S. Timmermans, Abductive analysis and grounded theory, *The SAGE handbook of current developments in grounded theory* (2019) 532–546.
- [58] K. Charmaz, I. I.-A. Abductive, Grounded theory coding, in: *Workshop presentation: The Changing Generations Project*. Dublin, Ireland, 2012.
- [59] M. Muller, Curiosity, creativity, and surprise as analytic tools: Grounded theory method, in: *Ways of Knowing in HCI*, Springer, 2014, pp. 25–48.
- [60] M. Lipscomb, Abductive reasoning and qualitative research, *Nursing Philosophy* 13 (2012) 244–256.
- [61] K. Rambaree, Abductive thematic network analysis (atna) using atlas-ti, in: *Innovative Research Methodologies in Management*, Springer, 2018, pp. 61–86.

- [62] J. Lindley, P. Coulton, Back to the future: 10 years of design fiction, in: Proceedings of the 2015 British HCI Conference, 2015, pp. 210–211.
- [63] M. Smyth, I. Helgason, Tangible possibilities—envisioning interactions in public space, *Digital Creativity* 24 (2013) 75–87.
- [64] M. Huusko, Y. Wu, V. Roto, Structuring and engaging: The roles of design fictions in a co-design workshop, in: Proceedings of the 30th Australian Conference on Computer-Human Interaction, 2018, pp. 234–241.
- [65] M. A. Blythe, P. C. Wright, Pastiche scenarios: Fiction as a resource for user centred design, *Interacting with computers* 18 (2006) 1139–1164.
- [66] M. Muller, T. Erickson, In the data kitchen: A review (a design fiction on data science), in: Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems, 2018, pp. 1–10.
- [67] R. Noortman, B. F. Schulte, P. Marshall, S. Bakker, A. L. Cox, Hawkeye-deploying a design fiction probe, in: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 2019, pp. 1–14.
- [68] M. Blythe, E. Encinas, The co-ordinates of design fiction: Extrapolation, irony, ambiguity and magic, in: Proceedings of the 19th international conference on supporting group work, 2016, pp. 345–354.
- [69] A. Dunne, F. Raby, *Speculative everything: design, fiction, and social dreaming*, MIT press, 2013.
- [70] M. L. J. Søndergaard, L. K. Hansen, Intimate futures: Staying with the trouble of digital personal assistants through design fiction, in: Proceedings of the 2018 Designing Interactive Systems Conference, 2018, pp. 869–880.
- [71] E. Encinas, M. Blythe, The solution printer: magic realist design fiction, in: Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems, 2016, pp. 387–396.
- [72] V. Fuchsberger, T. Meneweger, D. Wurhofer, M. Tscheligi, Apply now! fictional job postings as an instrument to discuss interactive futures of work, in: Proceedings of the 2017 Conference on Designing Interactive Systems, 2017, pp. 581–586.
- [73] C. Elsdén, D. Chatting, A. C. Durrant, A. Garbett, B. Nissen, J. Vines, D. S. Kirk, On speculative enactments, in: Proceedings of the 2017 CHI conference on human factors in computing systems, 2017, pp. 5386–5399.
- [74] B. Brown, A. S. Taylor, S. Izadi, A. Sellen, J. Jofish’Kaye, R. Eardley, Locating family values: A field trial of the whereabouts clock, in: International Conference on Ubiquitous Computing, Springer, 2007, pp. 354–371.
- [75] H. Hutchinson, W. Mackay, B. Westerlund, B. B. Bederson, A. Druin, C. Plaisant, M. Beaudouin-Lafon, S. Conversy, H. Evans, H. Hansen, et al., Technology probes: inspiring design for and with families, in: Proceedings of the SIGCHI conference on Human factors in computing systems, 2003, pp. 17–24.
- [76] B. F. Schulte, P. Marshall, A. L. Cox, Homes for life: a design fiction probe, in: Proceedings of the 9th nordic conference on human-computer interaction, 2016, pp. 1–10.
- [77] E. Cheon, N. M. Su, Futuristic autobiographies: Weaving participant narratives to elicit values around robots, in: Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction, 2018, pp. 388–397.
- [78] E. P. Baumer, M. Blythe, T. J. Tanenbaum, Evaluating design fiction: The right tool for the job, in: Proceedings of the 2020 ACM Designing Interactive Systems Conference, 2020, pp. 1901–1913.
- [79] J. Lindley, P. Coulton, Pushing the limits of design fiction: The case for fictional research papers, in: proceedings of the 2016 CHI conference on human factors in computing systems, 2016, pp. 4032–4043.
- [80] M. Muller, Q. V. Liao, Exploring ai ethics and values through participatory design fictions, *Human Computer Interaction Consortium* (2017).
- [81] M. Blythe, E. Buie, Chatbots of the gods: imaginary abstracts for techno-spirituality research, in: Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational, 2014, pp. 227–236.
- [82] M. Blythe, Research through design fiction: narrative in real and imaginary abstracts, in: Proceedings of the SIGCHI conference on human factors in computing systems, 2014, pp. 703–712.
- [83] A. H. Ambe, M. Brereton, A. Soro, L. Buys, P. Roe, The adventures of older authors: Exploring futures through co-design fictions, in: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 2019, pp. 1–16.
- [84] H. Draper, T. Sorell, Ethical values and social care robots for older people: an international qualitative study, *Ethics and Information Technology* 19 (2017) 49.
- [85] T. J. Tanenbaum, Design fictional interactions: why hci should care about stories, *interactions* 21 (2014) 22–23.
- [86] E. A. Buie, *Exploring techno-spirituality: Design strategies for transcendent user experiences*, University of Northumbria at Newcastle (United Kingdom), 2018.
- [87] W. A. Hamilton, N. Lupfer, A. Kerne, *Livedissent: A media platform for remote participation in activist demonstrations*, in: Proceedings of the 2018

- ACM Conference on Supporting Groupwork, 2018, pp. 257–266.
- [88] M. Wessel, A. Abdellatif, I. Wiese, T. Conte, E. Shihab, M. A. Gerosa, I. Steinmacher, Bots for pull requests: The good, the bad, and the promising (2022).
- [89] H. Candello, M. Pichiliani, M. Wessel, C. Pinhanez, M. Muller, Teaching robots to act and converse in physical spaces: participatory design fictions with museum guides, in: Proceedings of the Halfway to the Future Symposium 2019, 2019, pp. 1–4.
- [90] J. Auger, Speculative design: crafting the speculation, *Digital Creativity* 24 (2013) 11–35.
- [91] M. L. J. Søndergaard, L. K. Hansen, Periodshare: A bloody design fiction, in: Proceedings of the 9th Nordic Conference on Human-Computer Interaction, 2016, pp. 1–6.
- [92] B. Sterling, *Shaping things*, MIT Press, 2005.
- [93] M. Blythe, Research fiction: storytelling, plot and design, in: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, 2017, pp. 5400–5411.
- [94] M. Sturdee, P. Coulton, J. G. Lindley, M. Stead, H. Ali, A. Hudson-Smith, Design fiction: How to build a voight-kampff machine, in: Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems, 2016, pp. 375–386.
- [95] T. Jensen, P. Vistisen, Ethical design fiction: Between storytelling and world building, *The ORBIT Journal* 1 (2017) 1–14.
- [96] B. Sterling, Patently untrue: fleshy defibrillators and synchronised baseball are changing the future. *wired magazine*, 2015.
- [97] J. Lindley, A pragmatics framework for design fiction (2015).
- [98] F. Cafaro, L. Lyons, J. Roberts, J. Radinsky, The uncanny valley of embodied interaction design, in: Proceedings of the 2014 conference on Designing interactive systems, 2014, pp. 1075–1078.
- [99] B. Gaver, T. Dunne, E. Pacenti, Design: Cultural probes, *Interactions* 6 (1999) 21–29. URL: <https://doi.org/10.1145/291224.291235>. doi:10.1145/291224.291235.
- [100] K. Boehner, J. Vertesi, P. Sengers, P. Dourish, How hci interprets the probes, in: Proceedings of the SIGCHI conference on Human factors in computing systems, 2007, pp. 1077–1086.
- [101] W. Gaver, Cultural commentators: Non-native interpretations as resources for polyphonic assessment, *International journal of human-computer studies* 65 (2007) 292–305.
- [102] C. DiSalvo, Fej-142 spectacles and tropes: Speculative design and contemporary food cultures, *The Fibreculture Journal* (2012).
- [103] P. Coulton, J. G. Lindley, M. Sturdee, M. Stead, Design fiction as world building (2017).
- [104] S. P. Wyche, P. M. Aoki, R. E. Grinter, Re-placing faith: reconsidering the secular-religious use divide in the united states and kenya, in: Proceedings of the SIGCHI conference on human factors in computing systems, 2008, pp. 11–20.
- [105] D. Simons, Design for fairness in ai: Cooking a fair ai dish (2019).
- [106] P. Hayes, I. Van De Poel, M. Steen, Algorithms and values in justice and security, *AI & SOCIETY* (2020) 1–23.
- [107] A. Morrison, R. Tronstad, E. S. Martinussen, Design notes on a lonely drone, *Digital Creativity* 24 (2013) 46–59.
- [108] M. L. J. Søndergaard, Troubling design: A design program for designing with women’s health, *ACM Transactions on Computer-Human Interaction (TOCHI)* 27 (2020) 1–36.
- [109] L. Jönsson, A. Light, K. Lindström, Å. Ståhl, M. Tham, How can we come to care in and through design?, in: The 8th Bi-Annual Nordic Design Research Society Conference-Who Cares? 2-4th of June 2019 Finland, Nordic Design Research, 2019, pp. 1–8.
- [110] L. Devendorf, D. K. Rosner, Beyond hybrids: Metaphors and margins in design, in: Proceedings of the 2017 Conference on Designing Interactive Systems, 2017, pp. 995–1000.
- [111] S. Bardzell, J. Bardzell, Towards a feminist hci methodology: social science, feminism, and hci, in: Proceedings of the SIGCHI conference on human factors in computing systems, 2011, pp. 675–684.
- [112] D. J. Haraway, *Staying with the Trouble*, Duke University Press, 2016.
- [113] P. Thomas, M. Czerwinski, D. McDuff, N. Craswell, Theories of conversation for conversational ir, *ACM Transactions on Information Systems (TOIS)* 39 (2021) 1–23.
- [114] M. Chefitz, J. Austin-Breneman, N. Melville, Designing conversational interfaces to reduce dissonance, in: Proceedings of the 2018 ACM Conference Companion Publication on Designing Interactive Systems, 2018, pp. 219–223.
- [115] R. Louie, A. Coenen, C. Z. Huang, M. Terry, C. J. Cai, Novice-ai music co-creation via ai-steering tools for deep generative models, in: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, 2020, pp. 1–13.
- [116] D. Dablain, G. Siwo, N. Chawla, Generative ai design and exploration of nucleoside analogs (2021).
- [117] M. Farbood, H. Kaufman, K. Jennings, Composing with hyperscore: An intuitive interface for visualizing musical structure, in: ICMC, 2007.
- [118] A. Oeldorf-Hirsch, D. Gergle, ‘who knows what’

- audience targeting for question asking on facebook, *Proceedings of the ACM on Human-Computer Interaction* 4 (2020) 1–20.
- [119] N. S. Shami, Y. C. Yuan, D. Cosley, L. Xia, G. Gay, That’s what friends are for: facilitating who knows what across group boundaries, in: *Proceedings of the 2007 international ACM conference on Supporting group work*, 2007, pp. 379–382.
- [120] B. Biancardi, L. Maisonnave-Couterou, P. Renault, B. Ravenet, M. Mancini, G. Varni, The wonowa dataset: Investigating the transactive memory system in small group interactions, in: *Proceedings of the 2020 International Conference on Multimodal Interaction*, 2020, pp. 528–537.
- [121] S. Liang, Y. Luo, Z. Meng, Profiling users for question answering communities via flow-based constrained co-embedding model, *ACM Transactions on Information Systems (TOIS)* 40 (2021) 1–38.
- [122] C. Aragon, S. Guha, M. Kogan, M. Muller, G. Neff, *Human Centered Data Science: An Introduction*, MIT Press, 2022.
- [123] T. Ye, *Improving Worker Performance with Human-Centered Data Science*, Ph.D. thesis, 2021.
- [124] Y. Toyoda, G. Lucas, J. Gratch, The effects of autonomy and task meaning in algorithmic management of crowdwork, in: *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, 2020, pp. 1404–1412.
- [125] Z. Wang, C. S. Chong, L. Lan, Y. Yang, S. B. Ho, J. C. Tong, Fine-grained sentiment analysis of social media with emotion sensing, in: *2016 Future Technologies Conference (FTC)*, IEEE, 2016, pp. 1361–1364.
- [126] M. Skirpan, C. Fiesler, Ad empathy: A design fiction, in: *Proceedings of the 2018 ACM Conference on Supporting Groupwork*, 2018, pp. 267–273.
- [127] J. Anitha, Determinants of employee engagement and their impact on employee performance, *International journal of productivity and performance management* (2014).
- [128] S. Kular, M. Gatenby, C. Rees, E. Soane, K. Truss, Employee engagement: A literature review, 2008. URL: <http://eprints.kingston.ac.uk/4192/1/19wempen.pdf>, accessed: 2021-12-22.
- [129] W. H. Macey, B. Schneider, The meaning of employee engagement, *Industrial and organizational Psychology* 1 (2008) 3–30.
- [130] M. Agarwal, K. Talamadupula, F. Martinez, S. Houde, M. Muller, J. Richards, S. I. Ross, J. D. Weisz, Using document similarity methods to create parallel datasets for code translation, *arXiv preprint arXiv:2110.05423* (2021).
- [131] A. Veglis, T. A. Maniou, et al., Chatbots on the rise: A new narrative in journalism, *Stud. Media Commun* 7 (2019) 1–6.
- [132] Y. Zhou, Y. Koyama, M. Goto, T. Igarashi, Generative melody composition with human-in-the-loop bayesian optimization, *arXiv preprint arXiv:2010.03190* (2020).
- [133] T. Zhao, C. Chen, Y. Liu, X. Zhu, Guigan: Learning to generate gui designs using generative adversarial networks, in: *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, IEEE, 2021, pp. 748–760.
- [134] J. Ahmad, F. Zafar, Review of body area network technology & wireless medical monitoring, *International Journal of Information and Communication Technology* 2 (2012).
- [135] A. K. Teshome, B. Kibret, D. T. Lai, A review of implant communication technology in wlan: Progress and challenges, *IEEE reviews in biomedical engineering* 12 (2018) 88–99.
- [136] V. Lockton, R. S. Rosenberg, Rfid: The next serious threat to privacy, *Ethics and Information Technology* 7 (2005) 221–231.
- [137] U. Ahmad, H. Song, A. Bilal, S. Saleem, A. Ullah, Securing insulin pump system using deep learning and gesture recognition, in: *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/12th IEEE International Conference On Big Data Science And Engineering (TrustCom/Big-DataSE)*, IEEE, 2018, pp. 1716–1719.
- [138] H. Kassiri, N. Soltani, M. T. Salam, J. L. P. Velazquez, R. Genov, Battery-less modular responsive neurostimulator for prediction and abortion of epileptic seizures, in: *2016 IEEE International Symposium on Circuits and Systems (ISCAS)*, IEEE, 2016, pp. 1298–1301.
- [139] K. Abdelhalim, H. M. Jafari, L. Kokarovtseva, J. L. P. Velazquez, R. Genov, Neural synchrony-monitoring wireless brain implant for intractable epilepsy neuromodulation, in: *2013 6th International IEEE/EMBS Conference on Neural Engineering (NER)*, IEEE, 2013, pp. 65–68.
- [140] M. R. Rieback, B. Crispo, A. S. Tanenbaum, The evolution of rfid security, *IEEE Pervasive Computing* 5 (2006) 62–69.
- [141] S. J. Schlittmeier, J. Hellbrück, Background music as noise abatement in open-plan offices: A laboratory study on performance effects and subjective preferences, *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition* 23 (2009) 684–697.
- [142] A. B. Haake, Individual music listening in workplace settings: An exploratory survey of offices in the uk, *Musicae Scientiae* 15 (2011) 107–129.
- [143] R.-H. Huang, Y.-N. Shih, Effects of background music on concentration of workers, *Work* 38 (2011) 383–387.

- [144] Y. Furuya, K. Fukai, H. Furuya, M. Tatemichi, A pilot study on the effects of playing background music in the workplace on overtime work, *Environmental and Occupational Health Practice* (2021).
- [145] I. L. Ebert, Ai at work—promises and pitfalls of digital workplace monitoring, 2021. URL: <https://www.alexandria.unisg.ch/264999/1/SKEMA%20presentation%209%20Nov%20AI%20at%20work.pdf>, accessed 2 January 2022.
- [146] A. Aloisi, V. de Stefano, Essential jobs, remote work and digital surveillance: addressing the covid-19 pandemic panopticon, *International Labour Review* (2021).
- [147] B. Friedman, D. G. Hendry, Value sensitive design: Shaping technology with moral imagination, Mit Press, 2019.
- [148] L. V. Nägele, M. Ryöppy, D. Wilde, Pdf: participatory design fiction with vulnerable users, in: *Proceedings of the 10th Nordic Conference on Human-Computer Interaction*, 2018, pp. 819–831.
- [149] K. Baumann, B. Caldwell, F. Bar, B. Stokes, Participatory design fiction: community storytelling for speculative urban technologies, in: *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–1.
- [150] L. Forlano, A. Mathew, From design fiction to design friction: Speculative and participatory design of values-embedded urban technology, *Journal of Urban Technology* 21 (2014) 7–24.