

“That’s important, but...”: How Computer Science Researchers Anticipate Unintended Consequences of Their Research Innovations

Kimberly Do*†

Khoury College of Computer Sciences,
Northeastern University
Boston, Massachusetts, USA

Jiachen Jiang

Microsoft
Redmond, Washington, USA

Rock Yuren Pang*

Paul G. Allen School of Computer Science,
University of Washington
Seattle, Washington, USA

Katharina Reinecke

Paul G. Allen School of Computer Science,
University of Washington
Seattle, Washington, USA

ABSTRACT

Computer science research has led to many breakthrough innovations but has also been scrutinized for enabling technology that has negative, unintended consequences for society. Given the increasing discussions of ethics in the news and among researchers, we interviewed 20 researchers in various CS sub-disciplines to identify whether and how they consider potential unintended consequences of their research innovations. We show that considering unintended consequences is generally seen as important but rarely practiced. Principal barriers are a lack of formal process and strategy as well as the academic practice that prioritizes fast progress and publications. Drawing on these findings, we discuss approaches to support researchers in routinely considering unintended consequences, from bringing diverse perspectives through community participation to increasing incentives to investigate potential consequences. We intend for our work to pave the way for routine explorations of the societal implications of technological innovations before, during, and after the research process.

CCS CONCEPTS

• **Human-centered computing** → **HCI theory, concepts and models.**

KEYWORDS

Unintended Consequences, Computer Ethics

ACM Reference Format:

Kimberly Do*†, Rock Yuren Pang*, Jiachen Jiang, and Katharina Reinecke. 2023. “That’s important, but...”: How Computer Science Researchers Anticipate Unintended Consequences of Their Research Innovations. In *CHI '23, April 23–April 28, 2023, Hamburg, Germany*.

* Listed in alphabetical order. Both authors contributed equally to this research.

† Authored during REU at the University of Washington while affiliated with the Georgia Institute of Technology.

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

CHI '23, April 23–April 28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-XXXX-X/18/06.

<https://doi.org/10.1145/3544548.3581347>

Conference on Human Factors in Computing Systems (CHI '23), April 23–April 28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 16 pages.
<https://doi.org/10.1145/3544548.3581347>

1 INTRODUCTION

From smart glasses that invoke fears of surveillance [33] to chatbots that use racist language [72], our society must increasingly protect itself against the harmful effects of our own technological advancements — commonly referred to as unanticipated or unintended consequences [82, 95]. While industry is seen as the main offender due to its large user base and broad product impact, computer science research has experienced its own fair share of cases in which innovations have gone awry. For example, AI innovations enabling deepfakes have fostered the spread of disinformation [63]; deep learning models that predict a person’s sexual orientation have caused fierce backlash from the LGBTQ community [126]; language models have been shown to exacerbate inequalities [122], propagate social bias [68, 105], and intentionally produce discriminatory content [71]; and innovations in interaction design to improve usability have simultaneously widened disparities between the experiences of the demographic groups included or omitted from the research and development process [118].

Research in Science and Technology Studies (STS) has documented the hopes and challenges of technological utopianism for decades, but commonly focuses on industry and computer science practitioners, such as software engineers [65, 94]. Computer scientists in academia may face different and in some ways more complex challenges than practitioners when considering how to anticipate unintended consequences. For example, while their research can result in widely-adopted, non-commercial or commercial products, academic researchers commonly generate ideas and artifacts that are primarily used or extended by other researchers within academia [69].

The recent flurry of negative media about the adverse effects of technologies is spurring researchers in various computer science disciplines to more commonly examine ethical implications of their work. Human-Computer Interaction (HCI) researchers have explored the ethics of research and technology in workshops (e.g., [11, 127]) and various publications (e.g., [20, 22, 37, 48, 78, 80]).

Not too long ago, ethics was proposed as one of seven grand challenges for HCI [113]. The SIGCHI research ethics committee has been facilitating open conversations about ethical challenges in our communities through research ethics town halls and panels at CHI [39, 88], CSCW [24, 35], and GROUP [17]. Similar discussions for more ethical research have been pushed in disciplines such as natural language processing [56, 84], computer vision [29], virtual reality [8, 13, 100], robotics [59, 121], data management [3, 114], and data mining [50].

Most of the discussions on unintended consequences in computer science research have centered on issues that researchers face during the research process, such as when conducting user studies [39], using online data [123, 124], and crowdsourcing data collection [9]. What remains unknown is whether and how computer science researchers consider any potential unintended consequences of their innovations on society. How do they incorporate such considerations into their research practice, if at all, and what barriers do they encounter?

In this paper, we explored these questions through semi-structured interviews with 20 computer science researchers in various academic positions and sub-disciplines, including Accessibility, Augmented and Virtual Reality (AR/VR), CS Education, Computer Vision (CV), Fabrication, HCI, Machine Learning (ML), Natural Language Processing (NLP), Security, Social Computing, and Robotics. Our findings revealed researchers' current attitudes and practices surrounding the unintended consequences of their research innovations (Section 5.1). Concretely, we observed that researchers recognize the importance of this topic, but do not proactively anticipate potential unintended consequences in practice. In fact, thinking about possible unintended consequences of their research innovations is not an integral part of their research process. To further unpack the reasons for their (in)actions, we identified two main barriers to anticipating unintended consequences: (1) the lack of formal methods and guidelines to anticipate them (Section 5.2), and (2) academic practices that promote rapid progress and publications (Section 5.3).

Overall, this work presents an in-depth qualitative investigation of applied computer science researchers' current practices and challenges in dealing with unintended consequences. As a step towards more routinely considering these consequences in computer science research, we discuss directions for future research and provide actionable recommendations to the research community and individual researchers.

2 TERMINOLOGY

Already in 1936, Merton [82] coined the term *unanticipated consequences* to describe unforeseen, desirable or undesirable outcomes of one's action. Today, most researchers refer to unforeseen outcomes (the results of policies, technologies, or other "purposive social actions" [82]) as *unintended consequences*, though some have suggested that this term conflation has caused a loss of nuance [57, 95]. As we show in Figure 1, Merton's original term of *unanticipated consequences* suggests that such consequences are always unintended. In contrast, *unintended consequences* can be either unanticipated or anticipated. Parvin and Pollock have therefore argued that this lack of precision in terminology may lead

people "to abdicate responsibility for the perfectly foreseeable consequences of particular decisions" [95, p.323]. Hence, the use of the term *unintended consequences* may reduce accountability: "Phenomena described as unintended consequences are deemed too difficult, too out of scope, too out of reach, or too messy to have been dealt with at any point in time before they created problems for someone else. The descriptive approach works as a defensive and dismissive strategy" [95, p.322].

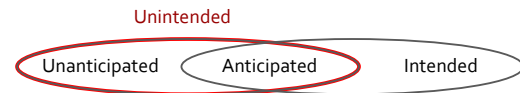


Figure 1: Terminology of anticipated, intended, unintended, and unanticipated consequences. Unintended consequences can be either unanticipated or anticipated.

In this paper, we use the term *unintended consequences (UCs)* to purposefully broaden the discussion to include both anticipated and unanticipated, positive or negative unintended side effects of technology on society. Our definition includes consequences that the instigators of an action (i.e., researchers and/or technology innovators) may not have addressed but could have foreseen. While these consequences can be positive or negative in nature (and oftentimes have different effects on a population), our work is inherently oriented towards considering negative UCs more than positive ones. In the remainder of this paper, we use "technology" for digital technology, such as hardware devices or software systems. We broadly refer to "society" at a regional, national, or international level.

3 RELATED WORK

Our work draws upon prior work studying values and ethics in digital technology [107] and is informed by discussions about the effects of digital technology on society in fields such as philosophy [62, 86], STS [64, 130, 131], social informatics [66], feminism [10, 49] and postcolonial theories [60]. We start by showing how researchers in these fields have long discussed various societal effects of technology before outlining the methods and approaches researchers and practitioners have developed for mitigating unintended consequences.

Critiques of Technology. Prior work in STS, and later in HCI, has provided critical analyses of the risks and benefits of technology in society since at least the 1960s [107, 116]. According to STS scholar Winner, technology "embodies specific forms of power and authority" [130] and technologists should "pay attention not only to the making of physical instruments and processes [...], but also the production of psychological, social, and political conditions as a part of any significant technical change" [131].

Work on the risks of technology has been published on a broad range of topics, including the Internet [70], health care information technologies [6, 54], mobile phones [87, 102], smart technologies [77], machine learning [25], and social media [30, 110, 111].

The examples above provide a critical lens of the role of technology in society and caution about the unknown and differential

effects on societies. Historically, however, most innovation research has focused on desirable and intended consequences [104, 116]. In 2009, Sveiby and colleagues suggested that this focus could potentially be due to a “pro-innovation bias among researchers and vested interests of funding agencies” [116]. Our literature review did not reveal whether this bias has changed in the years since.

However, we found many recent calls for more accountability for research innovations [40, 55, 83]. Several prominent computing conferences — such as the Conference on Neural Information Processing System (NeurIPS [15]), Annual Meetings of the Association for Computational Linguistics (ACL) [112], and the ACM Conference on Intelligent User Interfaces (IUI) [2] — have begun to experiment with ways to encourage or even require researchers to state both the positive and negative potential implications of their work in all paper submissions. Recent work has made several suggestions for such broader impact statements based on an analysis of these statements in NeurIPS conference proceedings [7, 76, 90]. After requiring all submissions to contain a section describing the impact of the work, NeurIPS has since transitioned towards a checklist system that offers additional guidance and adaptability [15].

The HCI community has been raising awareness of UCs of computing research through dedicated publication tracks (e.g., Critical Computing at CHI) and workshops [115, 120]. In particular, a CHI 2021 workshop explored how HCI researchers might think about and report potential negative consequences stemming from their research [115]. HCI researchers have also advocated for changes to the peer review process to reduce negative impacts of research innovations, suggesting that reviewers should routinely require that papers and proposals discuss potential adverse effects [55]. Given these calls for examining the societal impacts of technology, our work explores whether researchers adopt any methods for anticipating the UCs of their own work.

Anticipating and Mitigating Unintended Consequences of Technology. UCs are often dismissed as unavoidable because anticipating what may happen in the future can be hard [95] and uncertain [89]. However, HCI researchers have developed ethics-focused design methods to ensure the inclusion of various stakeholders in the design process (for an overview see [26]). One prominent example is the value-sensitive design (VSD) approach by Friedman, Kahn, and Borning [41], which can aid in understanding technology, its human value, and its context of use. The process aims to help product teams and researchers identify alternative approaches that better uphold their chosen values while accommodating the same constraints. A number of recent proposals have sought to bridge the gap between theory and implementation by creating toolkits meant for brainstorming about a product’s potential societal impacts. For example, the Envisioning Cards [42] present the VSD concepts in a clear and modular fashion [91]. In addition, stakeholder tokens also support a VSD stakeholder analysis [134]. Prior work in HCI has also recommended the use of Tarot Cards of Tech [81] and the Value Cards [106] for anticipating potential UCs of specific design choices.

Another approach for considering possible societal impacts is through design fiction [12, 18]. As a form of speculative design [31, 73], design fiction creates a fictional future world to think through sociotechnical issues that have relevance and implications for the

present [74, 133]. This practice has been used to reflect on potential downsides of public data [34], technology design [53], and research prototypes [109]. More recent work developed the design fiction memos method to explore how UX practitioners engage with ethical issues and social impact in their work [132].

The existing approaches to consider societal implications, however, were often assumed to be effective in practice [46] and might be difficult to evaluate [12]. While the toolkits often target designers and practitioners as users [46], applying them for research projects may pose additional complexities. We extend this line of work by inquiring into whether computer science researchers are aware of and proactively incorporate these tools in their research process. Our work also explores future design implications to support researchers to consider UCs in their research process.

Reacting to Unintended Consequences of Technology. Not much work has investigated how practitioners and researchers react to UCs in practice. Kling’s book on “Computerization and Controversy” shows how the power dynamics between programmers and their employers can prevent discussions of potential ethical issues in the products they work on [66]. As a result, computer science professionals may feel discouraged when reacting to potential or known UCs. Recently, an interview study showed that the Deepfake open source contributors felt unable to control downstream uses of their software, given the core principle of open source [128]. Researchers have occasionally written public posts in response to public backlash or negative press after deploying a research project [61, 92], but it is unclear whether they also do so when an incident is less public or when it has only been anticipated (but has not materialized). We fill this gap in prior work by studying whether and how academic computer scientists react if they discover that their work may have UCs.

4 METHODS

In this work, we conducted 20 semi-structured interviews to identify and understand whether and how computer science researchers from diverse sub-disciplines currently approach the potential UCs of their research innovations, what barriers they may encounter, and what design opportunities may exist to support this process.

Sampling and Participants. We used purposive sampling to select researchers in computer science who were affiliated with institutions in North America with very high research activities (R1). We focused on North American R1 institutions to reduce potential confounds related to the difference in academic culture and structures, such as funding applications and opportunities, requirements for promotion, and the structure of Ph.D. programs. Participants were required to work on applied research that has led, or could lead, to systems used by the general public.

We recruited participants via email after reading their webpages and publications to determine whether their research met our criteria and to ensure diversity in levels of research experience. The email briefly described the research goal of finding ways to support researchers in anticipating UCs. We recruited participants until we reached a sample that satisfied our goal of interviewing computer science researchers from diverse disciplines and seniority levels and until the interviews reached saturation. Although we attempted

Table 1: Overview of CS researchers in the study.

#	Institution	Position	Research Area	Released Public Products	Gender
P1	Public	PhD Student	NLP, HCI	No	Male
P2	Private	Assistant Professor	Social Computing	Yes	Male
P3	Public	Assistant Professor	ML	No	Male
P4	Private	Associate Professor	Security, Smart-phones, AI	Yes	Male
P5	Public	PhD Student	NLP	No	Male
P6	Public	Full Professor	NLP	Yes	Female
P7	Public	Full Professor	Robotics	Yes	Male
P8	Public	PhD Student	Computer Vision, ML	No	Male
P9	Private	PhD Student	AI	No	Male
P10	Private	PhD Student	AR, VR	No	Female
P11	Private	Assistant Professor	Brain-Computer Interfaces	No	Female
P12	Private	Postdoc	Accessibility	Yes	Female
P13	Private	PhD Student	Fabrication, Sensing	Yes	Male
P14	Private	Associate Professor	HCI	Yes	Female
P15	Public	Assistant Professor	AR, Accessibility	Yes	Male
P16	Public	PhD Student	CS Education, HCI	No	Female
P17	Public	PhD Student	CS Education	Yes	Male
P18	Public	Assistant Professor	AI, Robotics	No	Male
P19	Public	Assistant Professor	Security	Yes	Male
P20	Public	PhD Student	NLP	Yes	Female

to sample researchers from a variety of applied sub-disciplines, our findings may not encapsulate the thoughts and actions of all computer science researchers.

Our final sample included 20 computer science researchers (7 female, 13 male) from 10 different academic institutions across North America. All participants had built systems as part of their research, and 9 had released one or more systems as (part of) a public product. Participants held various academic positions in their respective computer science departments (see Table 1): 10 participants were Ph.D. students or postdocs, and the remaining 10 were assistant, associate, or full professors. Our participants worked in a variety of research areas: 9 participants described their work as being mainly in AI or related areas (e.g. CV, ML, NLP). The remaining participants worked on accessibility, AR/VR, CS education, hardware, social computing, robotics, and security, or a combination of the above. All participants had industry experience, meaning that they have either collaborated, interned, or obtained full-time positions in industry while working toward their research projects.

Interview Protocol. We prefaced our interviews by loosely defining "unintended consequences" as both desirable and undesirable outcomes of one's research, allowing participants to further elaborate on the term's meaning. We then divided our interviews into five sections: (1) participant research experience (e.g., research areas, educational and professional background); (2) prior experiences with UCs (e.g., from community norms in their sub-disciplines about considering UCs and/or their own research products have resulted in UCs); (3) understanding whether, when, and how they consider UCs in the research process; (4) understanding barriers to considering UCs in the research process; (5) understanding where researchers perceive opportunities to augment and improve the process to consider UCs. To avoid response bias, we started by explaining that the

topic of UCs is relatively new to computer science. In addition to their own experience, we asked questions about habits and norms in researchers' labs and communities to better understand the context for their opinions and avoid participants feeling accused or put on the spot. We used additional questions to probe three topics that came up repeatedly: understanding whether researchers actively anticipated UCs, understanding what may hinder them from doing so, and understanding the design opportunities to support them in anticipating UCs. See Supplementary Materials for the complete list of interview questions.

Nineteen interviews were conducted remotely over Zoom, and one was conducted in person. All participants consented to being audio-recorded. Each interview was between 30-45 minutes long. We made a financial donation of \$15 per participant to a COVID-19 relief fund to compensate each participant for their time. The study was determined as exempt by our Institutional Review Board (IRB).

Analysis. Our research team used an inductive thematic analysis process [21] where two researchers individually reviewed and conducted open coding on two interviews. Next, three researchers met to create and discuss the first draft of the codebook. Two members of the research team then independently coded several more interviews and refined or added them to the codebook, which was discussed with the full research team. Once consensus was reached, all interviews were re-coded using the revised codebook.

The final codebook contained 12 top-level codes relevant to how researchers have thought about, experienced, and responded to UCs; it also included codes related to attitudes towards UCs and support researchers needed to think about UCs. Finally, three researchers used affinity diagramming to develop themes based on our codes. Although we discussed both positive and negative UCs in our interviews, the nature of our research led us to focus on participants' reports of negative UCs. We slightly edited some of the quotes in Section 5 for readability.

Positionality. We acknowledge that our academic and professional backgrounds shape our perspectives on this topic. One author teaches computer ethics at an R1 institution. Collectively, we are US-based researchers at two R1 universities and a large US-based multinational corporation. Our academic backgrounds are in Computer Science, primarily as HCI researchers.

5 RESULTS

Our analysis surfaced three high-level themes. First, we describe current attitudes and practices surrounding the anticipation of UCs (Section 5.1). We then discuss how the lack of a formal method inhibits the anticipation and reaction to UCs by researchers (Section 5.2). Finally, we show how academic practices strain efforts to anticipate UCs (Section 5.3). Participants are identified with a "P". For a small number of sensitive quotes, the specific research products are omitted to provide an additional layer of anonymity.

5.1 Current Attitudes and Practices Surrounding UCs

In our interviews, 18 out of 20 participants explicitly mentioned that they had at least one prior experience where their research

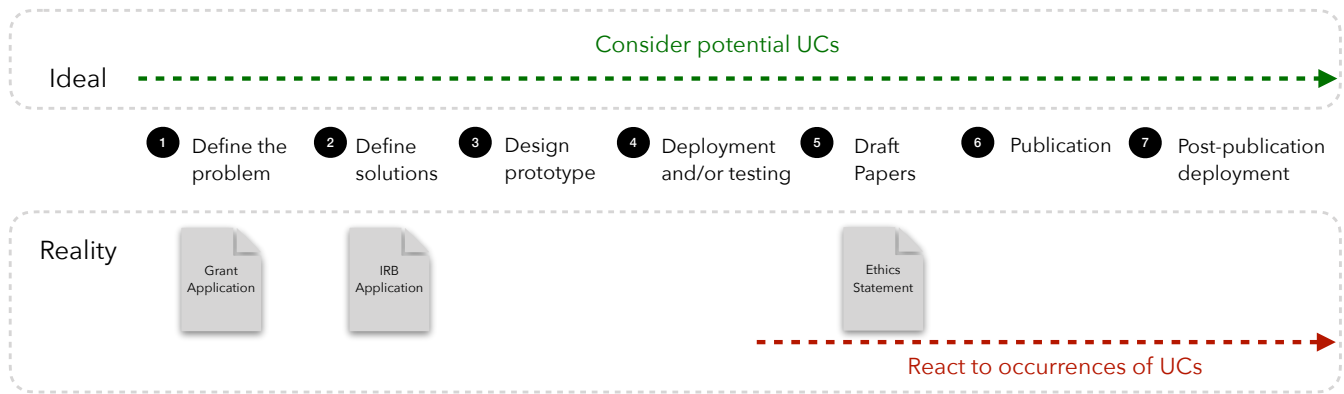


Figure 2: Stages in the research process during which UCs should be considered based on our interview results. Our findings suggest that considering potential UCs should ideally occur throughout the research process (i.e., the dashed green arrow), but that it is currently only done in reaction (i.e., the dashed red arrow) when writing grant applications, IRB applications, and ethics statements, if at all.

innovations had UCs after deployment or testing. All 20 participants unanimously emphasized the importance and responsibility of researchers to consider UCs throughout the research processes. To illustrate when our participants suggested ideally considering UCs, we present an overview of the research process in Figure 2. It shows that our participants indicated that UCs should ideally be considered throughout the research process, ranging from problem definition to publication and public deployment. In reality, we observed participants only anticipate UCs when they are required to do so, such as when writing broader impacts statements for grants, IRB applications, or ethics statements for conferences. We will provide more details describing how participants told us how they use these artifacts for reflecting on UCs in later sections (Section 5.1 and Section 5.2).

Despite the ideal circumstances for considering UCs, our analysis showed that none of our participants proactively consider them in the research process. Hence, **while participants see the importance of anticipating UCs, they rarely take actions to do so.** In fact, 10 participants self-described that they “do not spend enough time thinking about UCs,” and 4 participants explicitly mentioned that they “do not spend any time thinking about UCs.”

Despite not proactively anticipating UCs, our interviews revealed that **potential UCs are occasionally discussed through informal, serendipitous conversations, but there is no formal process.** Through our interviews, we found that researchers were sometimes made aware of potential UCs through informal conversations with their colleagues and collaborators. They recalled anticipating UCs during lab presentations or meetings, though infrequently. For example, P6 mentioned that collaborators sometimes share concerns, including potential UCs, on a case-by-case basis during project meetings. When asked when they think of UCs in their research, P6 remarked that they do so “on and off. I mean, the thought is generally there [...] because it’s not like the research projects change so often.” While giving informal research presentations in their lab, P1 welcomed feedback on potential UCs, but bewailed that “just talking about a research project doesn’t necessarily provide

an invitation for talking about UCs.” Many researchers shared similar experiences of learning about potential UCs through informal conversations or presentations, but also receiving limited feedback (P1, P4, P6, P7, P9, P10, P12, P16–17). For example, P16 shared her experiences on how “[Talking about UCs] is kind of something that just happens, during meetings, when if somebody just has a question to discuss, ‘Oh, I don’t know if that’s a great idea because of this and that.’” However, P16 also expressed that such an approach is “not like a formal [discussion to] make sure that we’re going through every aspect of this tool and making sure it’s not going to have these negative consequences.” Later P16 stated that the informal, or “accidental,” conversations on UCs are not an ideal solution.

In addition to these informal conversations, **participants commonly react to UCs only after creating tangible research artifacts which can range from deployment and/or testing to post-publication.** Of the 18 participants that shared specific experiences, 13 specifically recounted how experiencing UCs impacted later research decisions, such as making adjustments to a current research project (P5, P7, P17), identifying a new research direction (P12, P14, P19), terminating a research project or idea (P10–11, P13), or consulting expert assistance (P2, P4, P6, P11, P18). For example, P2 shared their experience when testing a content-sharing application in a user evaluation. P2 realized that some audiences found “some content might not be appropriate for them.” Therefore, they spent more time on content moderation than originally planned. Similarly, during post-publication, dissemination of results through methods such as social media may also heighten awareness that research innovations can have undesired societal impacts on researchers. For example, P10 told us:

“People put that [my] research project on Reddit. And there was like, a whole bunch of comments on Reddit on my research video [...] Like, there was something that I didn’t think about [...] And it was kind of something that actually led me to change research topics.”

P20 described their experience with a prototype system that received over 1 million adversarial examples from public auditing. They spent nights debugging the system, replying to users' critiques on Twitter, and issued a public disclaimer in the end. P20 recounted the outcome of their experience:

“When we initially released the archive paper [...], we didn't expect a lot of people would just play with a demo, because that's usually how these research prototypes are. [B]ut when we released the demo, suddenly, a lot of people played with it. So yeah, we never thought people would have had such an explosive discussion of this.”

As a result of reacting to UCs, participants became aware of the impact of their research products on users or society. While our participants took different actions to react to UCs, ranging from issuing public disclaimers (P6, P20), to debugging the system (P1-3, P5, P7-8, P14-16, P20), and shifting research direction (P2, P10), our interviews did not capture a standard procedure for anticipating UCs, such as a dedicated time and research procedure. Nonetheless, some participants regretted considering UCs only after a specific incident had occurred. As P5 put, *“A lot of people do treat [UCs] as an afterthought. And it's kind of unfortunate.”*

We also found that many **participants unintentionally deflect responsibility for considering UCs, believing that their research is unlikely to cause enough harm to warrant serious consideration of UCs**. Under the following subtheme, we describe researchers' acts of and justifications for this unintentional deflection.

Although all participants recognized the significance of anticipating UCs, some questioned whether researchers were personally responsible for anticipating the UCs of their research. While none of our participants stated that they ignore UCs altogether, several of them described that they are familiar with colleagues in the research community who disregard their work's societal implications (P6-8). As P8 put it, *“[Some colleagues might think] scientists are not responsible for broader impacts [...] They do research and then let policymakers or someone else think [about them].”* Similarly, P6 asserted that *“there are some machine learning people who think [...] I am going to focus on algorithms and not worry about social issues.”*

In contrast, some participants felt that researchers are responsible for the UCs of their work because they are the creators of new technologies. For example, P12 explained that:

“Researchers and anyone who is bringing these ideas to light, I think they need to be responsible for what they're saying, or at least be able to say whether they have discussed or kind of attempted to figure out what kind of impacts or implications are for whatever information they're presenting to the world.”

Despite the perceived responsibility of some participants, others worry that they lack the influence or foresight to engage with broader societal consequences. For example, P5 described working on a computer vision project that ultimately made it *“easier for the NSA [National Security Agency] to spy on people,”* but they neither anticipated nor reacted to this specific incident. They reconciled themselves with working on it because *“that is kind of why they were funding us.”* Similarly, P20 shared that it is *“extremely difficult to predict what can go wrong without seeing how [the research*

product] works in the real world, as an NLP researcher.” Moreover, some participants felt that how their research is used by others is beyond their own responsibility. To them, UCs were unfortunate, yet inevitable repercussions of research, but not necessarily a responsibility to consider or feasibly address.

We also found that the **Institutional Review Board (IRB) applications are mistakenly relied upon to alert researchers about the potential for UCs**. When asked about how they consider UCs in their research, some participants shared that they rely on the IRB application (P3, P8, P10-11, P13). P13 thought that *“the [IRB] committee does a good job and eliminates most of the foreseeable negative effects [of their submitted research projects].”* Likewise, P16 expressed similar confidence in the IRB application: *“The whole point of the IRB is that we're doing things ethically, right? [...] So yeah, I think [considering unintended consequences] falls under the same jurisdiction of [the IRB].”*

Others were well aware that the IRB process is not designed to anticipate UCs. For example, P10 noted that although the IRB considers *“[unintended] consequences in terms of individual participants, it does not specifically consider general consequences and consequently leaves the responsibility of considering unintended consequences to individual [researchers] to do it.”* P15 explained that the IRB application might not properly address non-human subjects research and stated that *“many AI projects where you don't interact with participants”* are not even considered by the IRB. Although the IRB application considers participant ethics, it is not designed to anticipate UCs on society beyond those related to participants. The observation made intuitive sense as the Common Rule, which governs the IRBs in the U.S., specifically disallows review of UCs to human society [38].

As noted above, all participants recognized the importance of considering UCs. Nevertheless, we observed that **participants underestimated the potential UCs of their own research because the primary goal of most research is to produce societal benefit**. In general, researchers equated good intentions with producing less social harm. For example, P4 explained that *“there is a pretty big distinction between [my] research [...], which tends to be very specifically targeted towards what hopefully will be societal good,”* and *“technologies that have caused more problems,”* specifying that their intent to create technologies with social benefit lessens the need to scrutinize potential UCs. P16 explicitly stated that *“researchers like us aren't trying to create something that could be harmful.”* Similarly, P6 noted that:

“By design, my research is really geared towards mitigating the problem that's out there, so it doesn't make sense for me to worry about the negative consequences [...] or think about the ethical concerns as much. Because it doesn't seem to apply.”

Although initially provided with a clarifying definition of UCs, multiple participants assumed that the positive, intended consequences of their research would reduce or eliminate its potential negative consequences, leading them to neglect considering UCs.

Additionally, other participants felt that there was no need for them to consider potential UCs because they perceived their research as unlikely to be misused. For example, P5 indicated that they are not as concerned with large-scale UCs because their work

is mostly hidden from the public. As an NLP researcher, they explained that “if anyone wanted to, like really generate fake news, at scale, they probably wouldn’t even have used our model.” Incidentally, many participants across multiple sub-disciplines shared this sentiment, but none could explain at what stage their research would require investigating potential UCs. Nearly all researchers explained that because their research created developmental technologies that might lead to “future work” [P20], they felt that any potential misuse of their research would not generate sufficient harm for them to need to consider UCs earlier on. In the words of P7, “because there are no consequences in academia, researchers have to get really worked up about [considering UCs].”

In summary, we found that our participants perceive considering UCs to be important, but they rarely take proactive steps to address this issue. Instead, they react to UCs when they occur. As P19 said, “I still think that’s important, but I just think it’s really hard.” Perhaps due to this attitude, many researchers unintentionally deflected responsibility for four key reasons: reliance on others to consider UCs, dependence on the IRB application to anticipate UCs, belief that research motives equated with reduced social harm, and doubt about the potential social impact of their work.

5.2 Researchers Lack Formal Methods and Guidelines for Anticipating UCs

Although participants understood the significance of anticipating and responding to UCs in their research, they generally felt unsure of how to approach this issue. Many participants reported that a lack of understanding and experience of UCs reduced their ability to anticipate and/or react to UCs. In this section, we delve into a list of barriers we identified from our interviews: a lack of systematic guidelines for anticipating UCs, a lack of experience and knowledge in UCs which hindered researchers’ ability to anticipate and react to UCs, and a lack of opportunities to work with collaborators from diverse backgrounds and skill sets which also hindered efforts to anticipate UCs.

Many researchers lamented a **lack of systematic guidelines for brainstorming about UCs and for knowing how to anticipate UCs**. P10 expressed that many conferences do not provide sufficient support — such as “predefined infrastructure or scaffolding approach for thinking about them”, despite requirements for broader impacts or ethics statements. As a result, researchers can feel that “it’s really on the individuals to do [anticipate UCs].” Compared to existing specific AI-related checklist [28, 45, 79], a guideline can inform “at what point you know if we could potentially have some very negative unintended consequences,” “who should we bring in,” “what kind of outside expert would be the most appropriate for this.” [P3]

Many participants yearned for guidelines to assist with anticipating UCs. P16 noted how a checklist for anticipating UCs could help researchers confirm “that they thought about [UCs] thoroughly.” In addition, P8, who expected to release an open-source CV demo, expressed a “need to develop some type of guidelines and policy for how much evaluation is sufficient [before deploying public technologies].”

Several participants also suggested the benefit of showcasing past UCs that others have anticipated or reacted. For example, P12 acknowledged how they had learned to anticipate UCs based on

previous experiences, so “a resource of common problems [...] of the past might be helpful, so I don’t make any of the mistakes that people have already made.” P10 suggested several potential resources for anticipating UCs including “some model examples of papers that have done a good job, researchers or research groups that have done a good job, [...] guidelines for how to approach [UCs] to begin with, [...] and thought experiments in thinking about how to how to come up with [UCs].” Similarly, P19 emphasized the impact of highlighting “very impactful paper[s] that include ethical statements and win an ACM award” as a way to not only provide guidance on how to structure writing about UCs, but also serve as “exemplars of papers” that demonstrate the importance of anticipating UCs to the academic community.

Similarly, **the broader impact or ethics statements — in grant applications and paper submissions — are insufficient in adequately accounting for UCs**. In grant applications, for example, P12 mentioned researchers might inaccurately portray the impact of their research, which “is typically framed in a positive light rather than a negative light,” in an effort to receive funding. The broader ethics statement that some conferences require were also discussed multiple times in our interviews. Although our participants viewed ethics and impact statements as the first step to anticipating UCs, they considered them as “superficial” solutions. P14 explained that, to many researchers, broader impact statements felt more like formalities that “researchers [have] to do to get publications, because that’s what helps them in their careers.” The broader ethics statement may make researchers only “think about unintended consequences or these other societal issues when they are writing”, rather than “when they are designing studies.” [P17] Additionally, we note that participants also feel conflicted about the effectiveness of these statements at guiding researchers to thoughtfully anticipate UCs throughout their research process. P12 worried that simply reporting UCs, rather than acting on them, does not sufficiently protect against the implications of negatively-impacting research. P8 added the limitation from a reviewer’s perspective:

“[A broader impacts statement] doesn’t really do anything, this is just like, a bandage on like a deeper wound of really investigating [...] whether or not a project should be pursued or not [...] Most reviewers when they read a paper to decide to accept to a conference or journal, won’t really seriously considered the broader impacts when deciding to accept or reject the paper.”

Moreover, many participants added that a **general lack of experience and knowledge in UCs hindered their ability to anticipate and react to UCs in their research**. Our interviews indicated that new faculty and junior researchers (e.g., graduate students, postdocs) were most likely to feel unsupported and ill-equipped when anticipating UCs. In particular, several participants shared their experiences with anticipating and addressing UCs as junior researchers. For example, P12 reflected on how, after experiencing a UC in their career, they realized that they did not have the research experience and foresight to anticipate any negative UCs, stating that they “[were] overly optimistic [...] about what the realistic [consequences were] going to be [and] sort of unaware of the [potential unintended consequences].” Other junior researchers also recognized how their limited research experiences prevent them

from anticipating UCs; therefore, to anticipate UCs, these junior researchers depend on their advisors to account for potential UCs (P1, P9–10, P13, P17). For example, P13 recounted an interaction with their advisor at the beginning of their Ph.D. where their *“PI basically turned [a research idea] down because it [could] have some negative use cases.”* Some junior researchers shared that they lack opportunities to consider UCs because of their mainly independent work. P1, a Ph.D. student, noted that their lack of collaborators led them to consult *“mostly [the] advisors in [their] labs”* to receive feedback about UCs. Additionally, P17 also acknowledged that *“maybe others have seen multiple things, or you get more perspectives or more voices. But that’s generally not the case for an average Ph.D. student, it’s usually a one-to-one relationship with an advisor when they’re starting [a Ph.D.]”*

In particular, junior researchers’ limited research experiences and mostly independent research work often meant that they were unable to properly account for UCs without the assistance of their advisors. Therefore, junior researchers might assume that more senior faculty would have already anticipated negative consequences. For example, P17, a Ph.D. student, explained that his advisor played a large role in screening his proposals for any potential UCs that they were simply unaware of. They explained that:

“I think [advisors] have realized [what the possible unintended consequences are] already [because] they went through the process, through writing what they have seen in the community [...] And if you spend enough time in the community, you get to know those norms implicitly.”

Further, several participants noted how new faculty also face difficulties when anticipating UCs. P18, a new faculty member, shared their thoughts on confronting UCs: *“I feel like there’s a responsibility there, but I also feel so incredibly ill-equipped to do anything useful about it.”* P19, another junior faculty member, also commented on the knowledge-based limitations that new faculty face in anticipating UCs and questioned their ability to consider *“the ethical reasoning about [a research project] if I’m not personally very educated on it.”*

Ironically, several faculty members noted that most students now enroll in ethics courses, and they therefore believed that students are sufficiently prepared to confront ethical problems in their research [P4, P6, P16]. However, our observation indicated that students and even junior faculty members may feel unable and unempowered to approach UCs without the support of their advisors and peers. We recognize how throughout the research pipeline, many researchers — including graduate students and faculty members — face knowledge barriers when anticipating UCs. Ultimately, these findings illustrate the lack of knowledge-based resources that all researchers face throughout the research pipeline.

Exacerbating these knowledge barriers, **sharing and learning from prior experiences in computer science is largely unsupported, both individually and at a research community level.** Participants stated that reporting UCs during or after publication made it difficult for other researchers to gain awareness of potential pitfalls others have encountered. For example, P11 shared their process for anticipating UCs and explained how they reviewed *“other papers that do similar work to see how they’ve done if they documented*

anything, [but] a lot of times they don’t.” Similarly, P19 suggested *“having checkpoints where you share sort of interim progress that you make with the community”* so that other researchers can identify and anticipate when potential UCs could occur during the research process. Notably, P14 and P16, who realized the UCs of their work only after publication and deployment, reported that they could not revise their papers without either retracting the work altogether or adding significantly more contributions to the original work. As P16 explained:

“And now [that] we’re seeing that there are some unintended consequences, there should be a way for me as the researcher, so that I have the responsibility to go over to that, wherever it’s published, and comment on it and be like, here’s an update or whatever, so that the community knows.”

To these researchers, reporting UCs was burdensome and potentially harmful to their careers, but they both felt willing and obligated to report these updates. Nonetheless, due to significant structural barriers, none of these researchers were able to share these contributions.

Lastly, many researchers suggested that **a lack of opportunities to work with collaborators from diverse backgrounds and skill sets also hindered efforts to anticipate UCs.** In general, many researchers believed that more opportunities to work with collaborators from diverse backgrounds and skill sets could better support anticipating UCs by enabling the exchange of different viewpoints (P1–2, P8, P11–12, P14–15). For example, P12 asserted the importance of *“be[ing] able to get feedback from people with different life experiences and expertise.”* In contrast, several researchers reported attending ethics workshops at conferences enabled them to gain feedback about potential UCs from others (P1, P4, P17). For example, P9 specified that meeting with others at conference workshops allowed *“leading researchers and also students [to] come together, [and] would help the community a lot in terms of figuring out [broader impacts].”* Additionally, two researchers mentioned their participation in university-level ethics board (P12, P18). In particular, P12 shared how their experience with volunteering at their university’s ethics board enabled them to *“identify problem[s] that a lot of people face...and provide a resource [of] people who had different research experience in diverse areas”* P1 also expressed their enthusiasm for greater diversity in computer science labs:

“I would [want] to have more people like even just people in my lab that have more diverse experiences and backgrounds. I think that we already do have that and that’s why I’m saying that I’ve noticed that this is super helpful when talking about unintended consequences.”

Further, participants explained how several barriers (related to current academic practices, Section 5.3) prevented them from working with others. For example, P3 shared how they usually had enough time to ask for external help only to conduct “sanity checks” but otherwise lacked time to officially “bring in outside experts.” Other participants intentionally chose to collaborate within their own labs because it was assumed that their peers would offer sufficient feedback and suggest potential UCs during lab meetings, P13 noted that their lab did not seek external collaborations.

Participants also specified that diversity and inclusivity had directly impacted their research projects. In particular, P15, an accessibility researcher, shared how a recent collaboration “with a researcher with a disability” helped their lab “[learn] a lot from their experience and [brought] in perspective from those communities.” Additionally, P14 shared that lacking a diverse team of researchers on a past research experience led to a UC that users found “deeply offensive”; P14 also highlighted how this experience led to future collaborations with more diverse teams and the use of participatory design methods. Generally, researchers agreed that diverse researchers could better help anticipate UCs by considering a greater variety of perspectives.

5.3 Academic Practices Strain Efforts to Anticipate UCs

Despite growing attempts to encourage anticipating UCs in research, we found that the **“move-fast” academic practice strains efforts to consider UCs**. Under this theme, we identified a list of barriers to elaborate how existing academic practices influence researchers’ actions and attitudes towards UCs. In particular, we describe how academic pressures to publish frequently impacts researchers’ considerations of UCs.

First, several participants described that they did not have enough time to devote to thinking about UCs. Participants felt that **considering UCs was additional work that might conflict with their goals to publish quickly**. To fully consider the social impacts, participants need to balance between moving fast and slowing down to brainstorm the future. Usually researchers chose to move fast, despite their hope to slow down ethically. P3 described their attitudes towards considering UCs while working on research projects: “I think we all have a tendency to try to move fast and [think] ‘Oh no, I don’t want to stop and think about [unintended consequences], I want to keep going [...] to get the results.’” For example, P1 described that they prioritized the “technical or theoretical hurdles” and considering UCs is “never one of those hurdles [...] never something that you explicitly approached with.”

Participants frequently commented that this academic pressure created incentives for producing more publications, rather than promoting ethical considerations. To some participants like P8, the result of this incentives system is the widespread belief that UCs are non-essential and “just adds more burden on the researcher.” P19 explained: “In academia, the incentives are lacking [and] misaligned, which kind of feeds into that broader system and whole obsession over count metrics, so people try as hard as possible to publish as much as fast as possible.” Similarly, P2 commented: “Incentives within [academic] systems may lead [researchers] to prioritize or de-prioritize things based on what’s most valuable: so in academia that will be publications.” P14, an associate professor in HCI, also scrutinized the research practices of AI and ML researchers, who “can start a project at the first one month and have a written paper in three months.” They continued that this oftentimes leads to “unethical work that have high citation rates and high rates of turning out publications.”

The **“move-fast” academic practice makes junior researchers difficult to consider UCs**. P7, a full professor at an R1 university, elaborated:

“[J]unior faculty, sometimes they don’t have the time, the mind space, or the clarity to think about these issues, because [they’re] just busy doing other things. I don’t fault them for that. It’s just [I] don’t think we’re creating an environment such that they have enough oxygen to think about these issues. So I think that’s something for us as a community to think about is just how do we, how do we build a space? And how do we build the time such that people are educated about these issues, and then have the opportunity to develop nuanced thoughts on them?”

Ultimately, these comments by our participants revealed that the unfortunate by-product of the academic pressure to publish inadvertently encouraged a system where researchers may fail to have the time for, or even see the value of, considering UCs in their research.

On a positive note, while participants expressed frustration with the academic system, some also mentioned how things might be slowly starting to change. P15 shared an experience where they worked with a student who pointed out significant ethical considerations that they had overlooked for a research project using information learned from “a particular [ethics] class that that [the student] took”.

As P6, a full professor in NLP, told us:

“I have more hope for the new generation, the new generation of students, because they tend to be more concerned about it in general, which might be to do with the fact that they somehow learned about it during their college education, for example, they may have had a class about ethics and inequality.”

Hence, with an increase in ethics courses, a general increase in awareness, and with systemic changes that mitigate barriers for considering UCs, academics may become more conscientious and incentivized to think about societal impacts of their work in the future.

6 DISCUSSION

Our findings demonstrate that computer science researchers in our study do not formally consider any potential societal impacts of their research innovations, despite perceiving it as important. Participants contemplated potential societal impacts of their innovations only in hindsight, such as when they were required to write a broader impact statement for a conference. They responded to individual incidents, such as when receiving participant feedback or getting bad press, instead of actively considering the topic. We contend that this observation is troublesome, as argued by Pillai et al. [44, p.2]: “Unless ethics is integrated in every aspect of the design process and educational curriculum, it is bound to be an afterthought and thus inadequate in identifying and addressing ethical issues.”

While our participants generally suggested the need to proactively consider UCs, we identified various knowledge and structural barriers that currently prevent them from doing so. First, participants felt that the lack of considering UCs is due to academic practices promoting fast progress and publications. In their eyes, the pressure to publish encourages researchers to de-prioritize and

Table 2: Five Causes of Unanticipated Consequences by Robert Merton [82]

Ignorance: Lack of knowledge, experience, expertise, and prudent investigation of a problem.
Errors: Incorrect reasoning, analysis techniques, and interpretation of a problem.
Imperious immediacy of interest: Actor’s paramount concern with the foreseen immediate consequences excludes the consideration of further or other consequences of the same act.
Basic Values: No consideration of further consequences because of the felt necessity of certain actions enjoined by certain fundamental values.
Self-defeating prophecies: Predictions are frequently not sustained precisely because the prediction has become a new element in the concrete situation.

resist any form of stepping back and considering long-term effects. Several interviewees felt that spending time evaluating potential UCs slowed their momentum when conducting actual research.

Second, we identified that researchers lack guidelines, tools, and methods for considering UCs. None of our participants reported using any of the existing tools for brainstorming about potential societal impacts, such as the Envisioning Cards [42] or the Tarot Cards of Tech [5]. They also suggested that there are no approaches or processes for thinking through the societal implications of their research that they felt genuinely satisfied with.

Third, participants mentioned the lack of demographic diversity in collaborators and other academics, which they felt could enrich different viewpoints and experiences. This is a known structural issue in computer science, an occupational group that is heavily skewed towards white males [32]. Although the U.S. student population is becoming more diverse, faculty members and general academics remain predominantly white and male [19]. Furthermore, in a number of fields, most research published at high-profile venues is conducted in Western countries, with only Western participants [75]. This means that research and innovation processes are distanced from the lives and experiences of many of their future users. Additionally, innovations can affect different kinds of people in unpredictable ways. This lack of diverse viewpoints and experiences characterized innovations that are biased against minority populations [52, 97].

Our exploration of the barriers to researchers’ (in)actions resonates with Merton’s five causes of UCs (see Table 2) [82] and enables us to place his theory in the academic context today. For example, Merton described how ignorance (i.e., a lack of knowledge and experience) can lead to UCs. We saw ignorance being part of the issue in the results of our study, with many researchers lacking the know-how to anticipate UCs and the experience of thinking about potential impacts from diverse perspectives. Merton also suggested that UCs can be caused by an “imperious immediacy of interest,” with people being driven to focus on the foreseen, desired consequences. As our participants mentioned, academic pressures to publish could exacerbate the desire to meet *intended* consequences, such as developing a technology for the purpose of publication and/or getting a degree or promotion. They may therefore unconsciously or consciously ignore unintended effects.

Our results indicate that researchers would more routinely consider UCs if they felt that their research outcomes had a future impact on society. As it is, many deflect responsibility to those whose research they believe is more likely to influence products.

The insight surprised us given that all interviewees were selected based on their work in applied research areas, and all had previously worked for research spin-offs, open-source projects, or companies. It suggests that there may be a gap between the impact academics think their work will have and the actual risk of UCs. However, the finding is in line with recent work [7, 90] analyzing the text included in the broader impact statements for NeurIPS conference papers. Their results show that authors rarely state who is responsible for preventing negative impacts, and if they do, it is often a call for action rather than a statement of adopting personal responsibility.

Several participants mentioned that they relied on their IRBs to alert them to potential UCs, wrongfully assuming that this “falls under the same jurisdiction.” While providing such structure or access could certainly help, it may risk replacing researchers’ responsibilities with “legalistic bureaucracy,” a known criticism of the IRB [22]. This is the case for ethics checklists as well, such as those designed to guide practitioners’ development of AI systems, which have been found to obfuscate responsibility while additionally being too abstract or ignored [79]. Nevertheless, the finding emphasizes the need for computer science researchers to receive more guidance in thinking through societal implications, be it through an IRB-like structure or access to ethic experts who can provide guidance.

The deflection of responsibility could be attributed to insufficient prior experiences with computing ethics and awareness of cautionary tales of innovations that had negative societal impacts. In fact, many participants cited their limited experience as a barrier to thinking about potential ethics pitfalls. This may be slowly changing as more universities now offer ethics classes as part of their computer science curriculum [36], which, as several senior faculty interviewees noted, is how their students gain experience with considering these consequences. With an increase in ethics education and a generally heightened mindfulness of these issues due to media attention and public outcries, we may look back and be surprised that computer science researchers do not regularly consider ethical consequences in the research process. We remain hopeful, but our findings also show that overcoming barriers will require a holistic approach to restructuring academia. In the following, we discuss how to better support researchers anticipating UCs.

7 SUPPORTING CONSIDERATION OF POTENTIAL UCS OF RESEARCH INNOVATIONS

Our analysis suggests that there is still a long way to go before UCs are fully taken into account by researchers, despite recent calls to do so. Our work, however, does not contend that all computer science researchers should be ethicists to address the issue. Instead, by properly anticipating these incidents of varying severity [47], researchers might also avoid suffering from their own reputation tarnishing and unexpectedly developing harmful technologies. Releasing technology innovations into the wild without fully considering their effects on society is arguably a very large human-subject experiment with unknown outcomes. Our work reveals that considering societal impacts, at this point, is primarily an individual responsibility and that some form of oversight should be in place. While structural changes may be needed to overcome

these issues, we offer concrete suggestions for tangible next steps as follows.

Collect and disseminate case studies of UCs that were the result of computer science research innovations. Our findings suggest that participants sometimes felt their own research to be too prototypical and too removed from public release to cause UCs. Others believed that certain research, especially if the goal is to address a societal issue, was less likely to result in UCs. These findings suggest the need for information that could raise awareness of UCs that resulted from varied research innovations. Concretely, we suggest collecting and disseminating case studies of the societal impacts of computer science research, which can serve as an informational resource for teaching new and experienced researchers about prior work that has had (differential) negative effects on society. It is crucial that such case studies and other relevant resources demonstrate different aspects of UCs on innovations that might be otherwise overlooked. Moreover, collecting such resources should be an invitation to collectively learn from past mistakes instead of finger-pointing, and their submission should be encouraged, such as by providing specific awards.

Develop tools that support learning and brainstorming about UCs. Our study shows that researchers need more structured guidance to think through potential societal impacts. While some tools already exist, such as the Tarot Cards of Tech or Envisioning Cards, these often target practitioners and are not always suitable for research artifacts. They also require much time and cognitive effort to brainstorm about UCs. Moreover, anticipating the downstream uses and effects of research innovations is a notoriously difficult problem [98]. Since a majority of our participants relied on their peers’ knowledge of prior UC experiences, creativity support tools, such as those that cluster ideas from previous users to support the generation of more diverse ideas [108], could promote sharing of UCs within communities. Such collaborative ideation tools could also serve as one approach for engaging citizen scientists in the research process.

Increase access to input from diverse people in research design and development. Our results show the benefits of more demographic diversity among researchers, so their projects can be conceived and evolve with input from people with varying values and experiences. In particular, several participants mentioned how, especially when attending alongside a diverse panel of researchers, participating in ethics workshops at conferences supported conversations and considerations about ethics in their research. This suggestion also reflects discussions by Kling [65], who encouraged computer science practitioners to apply diverse perspectives from other fields (e.g., the social sciences) to recognize potential societal implications, and Hankerson et al. [52], who suggested that hiring more diverse people at every level could help mitigate racial biases in technology.

While these suggestions are long-term goals, a more immediate suggestion, inspired by human-centered design practices, is to engage with diverse participants of varying backgrounds, skills, and characteristics to ideate and test new technologies. For example, researchers have found that unrepresentative participant sampling

may lead to the creation of non-generalizable technologies that amplify existing inequities [75, 103]. Similarly, to prevent further social inequity in technology, we also propose the use of participatory and co-design methods, which have been shown to more equitably and effectively inform technologies by directly identifying user needs [117, 125].

Second, we suggest inviting the public to routinely shape research ideas and ongoing projects with comments and feedback. For example, Johnson and Crivellaro designed a community panel to create dialogic spaces that foster critical engagements with technologies and social issues for the purpose of reviewing research proposals on HCI [43]. Many citizen science projects witnessed radical improvements due to public contributions [51]. These projects offer a give and take, with researchers receiving help collecting, annotating, and analyzing data in exchange for an educational gain for citizen scientists [85, 101]. A similar involvement of citizen scientists in evaluating the societal impacts of research innovations (at all stages of the research process) could benefit both researchers and the public.

Of course, such involvement requires precautions. For example, participatory methods might bring unequal power dynamics and risk over-reliance on citizen scientists who may not be experienced and knowledgeable in examining these issues [16]. Researchers may also be hesitant to openly discuss their ideas before they are published. This may be mitigated by allowing the preregistration of ideas (similar to what has been encouraged for research studies [27, 93]). Overall, we believe that allowing the public to have a voice in evaluating the societal implications of research projects on them will be a challenging but rewarding endeavor.

Increase incentives to investigate UCs at various stages of the research process. Our findings suggest that academics may have a lower incentive to investigate UCs than companies do, because they often do not perceive their research as immediately affecting large numbers of users and because they are less likely to face direct public backlash. While any additional task will be perceived as burdensome, we can learn from several efforts to improve scientific practices. For example, conferences have included tracks that auto-accept papers if researchers preregister their studies [58]. Several venues have also encouraged the replication of studies to improve the quality of science [129]. Similarly, computer science venues could encourage the publication of auditing research artifacts [99] by providing special tracks. To support such early audits, these venues could additionally allow for submissions of experience reports that describe how a research project was modified or even canceled due to the discovery of unanticipated societal impacts. According to our participants, researchers would greatly benefit from learning about such experiences to prevent UCs in the future.

Additionally, conferences, research methods classes, and advisors should reframe finding negative UCs as an opportunity. Oftentimes, analyzing a project for disparate outcomes can “serve as a starting point for another research project” [2]. For example, prior work on language models has suggested that discussing the potential harms of this technology can “stimulate efforts to study and mitigate them” [23, p. 34]. Similarly, research communities should consider rewarding researchers for uncovering and addressing societal consequences and provide venues for openly debating

whether research efforts should continue, be abandoned, or change direction.

Finally, as our participants suggested, adding “lessons learned” to existing research publications should be made possible in digital libraries to allow others to learn from these experiences. Collecting and sharing these experiences with researchers may also help overcome the common “not me” attitude that we have seen in our study by showing researchers that no matter how well intentioned, almost all research can impose negative consequences that must be considered.

Front-load considerations of societal impacts. Recently, conferences require or nudge authors to consider UCs, e.g., by including broader impact statements in papers. However, critics have raised concerns that these statements contain speculative fiction and are published too late in the research process [4]. Front-loading considerations of societal impacts in the research process can avoid the point of no return, as some of our participants described it. One possibility is encouraging pilot studies to confirm ethical practices before conducting large-scale human-subjects experiments [119]. Another possibility is to require submitting an analysis of potential UCs to funding agencies along with research proposals. Funding agencies commonly request broader impacts statements, but these are mostly used to advertise the positive impacts of a research project rather than to reflect on potential negative consequences.

Two challenges must be addressed to make this happen: First, authors of a grant proposal would need adequate guidelines and tools to think through societal impacts. Second, funding agencies and their reviewers would need to learn how to evaluate such societal impact statements. A similar process was successfully tested at an academic institution [14], though it was cautioned that it can be perceived as “burden” without additional scaffolding.

Therefore, one immediate next step is to invest in research efforts that can provide guidance to authors and reviewers. Institutions could further assist researchers with regular consultation by hiring an ethics advisor similar to the ethics consultation service in most medical programs [1, 96]. With the onset of more advanced technology applications, having a technology ethics advisor on staff to answer researcher questions and review proposals and papers might alleviate burdens for researchers on their own.

Provide guidelines for reacting to societal impacts. In addition to the lack of guidelines for *considering* UCs, our participants frequently mentioned the lack of guidelines for *reacting* to them once they were discovered. Their hesitation to consider societal impacts may come from knowledge barriers regarding anticipating and reacting to UCs. In line with Knight’s recommendation to addressing UCs, we offer two suggestions that aim to “increase knowledge” about UCs and combine uncertainties within a large-scale community [67]. First, we advocate creating opportunities where researchers with similar research topics can share their experiences with considering and addressing UCs. These opportunities, which could be digital platforms (e.g., online forums) or physical venues (e.g., conferences and workshops), would not only support researchers by recognizing common UCs within their sub-disciplines but would also create community and dialogue around UCs. Second, funding agencies could provide resources similar to

the IUI 2022 program chairs, who have put together a list of links to papers and tools for authors to audit their projects [2]. Providing them with checklists and tools to evaluate these analysis reports will be essential to avoid ethics washing [98]. Altogether, providing guidelines for such decisions is desirable; we believe that HCI researchers, given the field’s interdisciplinary nature and focus on societal impact, are uniquely equipped to drive research in this direction.

8 LIMITATIONS AND FUTURE WORK

Our work provides insight into the barriers that prevent computer science researchers from considering and addressing UCs of their work. One limitation of our work is the relatively limited sample size and diversity of participants in this study. We selectively spoke to North American computer science researchers from R1 research universities. Because academia can be structured differently across countries, our findings may not reflect either the resources or challenges that academics in other parts of the world encounter when considering societal impacts of their work. The focus on R1 institutions additionally means that we cannot conclude the habits and challenges in considering UCs are similar to R2 or special focus institutions. Moreover, because we employed purposive sampling, our findings may not generalize to all computer science researchers; in particular, we suspect there may be differences if researchers come from marginalized or minority backgrounds, if their research addresses ethics, or if their research is further removed from applications. Future work is needed to investigate whether our findings and suggestions for supporting computer science researchers in anticipating UCs generalize to other research institutions within the US and to other countries.

An additional limitation is our focus on academic computer science researchers whose applied research products have led to systems used by the general public. This might have overlooked opinions by researchers in industry who may be subject to different policies and organizational structures within their companies. We might have missed opinions by researchers in other fields that are deeply affected by UCs in technology. Future work should explore how computer science researchers in their specialized sub-disciplines with varying demographic backgrounds, work experiences, and research experiences enrich our findings.

Our results are also impacted by the possibility of response bias, a common issue in interview studies. Given the heightened awareness of societal implications and the blame associated with it, the risk is that participants may have appeared more concerned about UCs than they actually are and may have downplayed any UCs that they have experienced themselves. We did not perceive this as an issue in their responses, but our findings on the perceived importance of proactively considering UCs should still be taken with a grain of salt. Future work could build on these findings with an anonymous survey or with longitudinal studies on people’s motivations and actions for considering UCs.

Additionally, our work mainly focuses on how researchers anticipate unintended consequences. Future work should explore the opportunities and challenges that researchers encounter when reacting to unintended consequences in greater detail. Based on our initial findings from this work, we assume that researchers face

similar structural and knowledge barriers when responding to unintended consequences during the research process.

9 CONCLUSION

In this paper, we share how computer science researchers anticipate the UCs of their research innovations and the challenges they may face from doing so throughout the research process. Through our interviews with 20 computer science researchers affiliated with North American research universities, we learned that our interviewees do not have a specific process or strategy for considering UCs. Our interviews surfaced two major barriers to anticipating UCs for researchers including a lack of formal methods and guidelines for anticipating them and academic practices that promote fast progress and frequent publications. Based on our findings, we outline key opportunities to support researchers in these efforts by incentivizing the investigation of UCs throughout the research process, the creation of new tools to support brainstorming, and the implementation of reactionary guidelines. We hope that our work will encourage further discussions on how academic innovators can be supported in the prevention of unintended effects of technology on society.

ACKNOWLEDGMENTS

We thank our participants and the anonymous reviewers for their valuable feedback and suggestions. We also thank Sandy Kaplan and Mary Peng for their help revising this paper. This work was funded by the National Science Foundation under award IIS-2006104. Any opinions, findings, conclusions, or recommendations expressed in our work are those of the authors and do not necessarily reflect those of the supporter.

REFERENCES

- [1] American Medical Association 2021. *Ethics Consultations*. American Medical Association. Retrieved 2021-09-14 from <https://www.ama-assn.org/delivering-care/ethics/ethics-consultations>
- [2] IUI Program Chairs 2022. 2021. Reflecting on Societal Implications of IUI Research. https://iui.acm.org/2022/societal_impact.html, last accessed: September 6, 2021.
- [3] Serge Abiteboul and Julia Stoyanovich. 2019. Transparency, Fairness, Data Protection, Neutrality: Data Management Challenges in the Face of New Regulation. *J. Data and Information Quality* 11, 3, Article 15 (jun 2019), 9 pages. <https://doi.org/10.1145/3310231>
- [4] Eytan Adar. 2018. Negative Potential in CS. <https://medium.com/@eytanadar/negative-potential-in-cs-b26964b4f502>, last accessed: September 6, 2021.
- [5] Artefact Group. 2017. The Tarot Cards of Tech. <http://tarotcardsoftech.artefactgroup.com/>, last accessed: September 6, 2021.
- [6] Joan S Ash, Dean F Sittig, Richard H Dykstra, Kenneth Guappone, James D Carpenter, and Veena Seshadri. 2007. Categorizing the unintended sociotechnical consequences of computerized provider order entry. *International journal of medical informatics* 76 (2007), S21–S27.
- [7] Carolyn Ashurst, Emmie Hine, Paul Sedille, and Alexis Carlier. 2022. AI Ethics Statements: Analysis and Lessons Learnt from NeurIPS Broader Impact Statements. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 2047–2056. <https://doi.org/10.1145/3531146.3533780>
- [8] Jonathan Barbara, Hartmut Koehnitz, and Ágnes Karolina Bakk. 2021. The Ethics of Virtual Reality Interactive Digital Narratives in Cultural Heritage. In *Interactive Storytelling: 14th International Conference on Interactive Digital Storytelling, ICIDS 2021, Tallinn, Estonia, December 7–10, 2021, Proceedings* (Tallinn, Estonia). Springer-Verlag, Berlin, Heidelberg, 288–292. https://doi.org/10.1007/978-3-030-92300-6_27
- [9] Natã M. Barbosa and Monchu Chen. 2019. Rehumanized Crowdsourcing: A Labeling Framework Addressing Bias and Ethics in Machine Learning. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*

- (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300773>
- [10] Shaowen Bardzell. 2010. Feminist HCI: Taking Stock and Outlining an Agenda for Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '10). Association for Computing Machinery, New York, NY, USA, 1301–1310. <https://doi.org/10.1145/1753326.1753521>
- [11] Oliver Bates, Kathy New, Samantha Mitchell-Finnigan, Matthew Louis Mauriello, Christian Remy, Roy Bendor, Samuel Mann, Simran Chopra, Adrian K. Clear, and Chris Preist. 2019. Towards a Responsible Innovation Agenda for HCI. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI EA '19). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3290607.3299017>
- [12] Eric P.S. Baumer, Timothy Berrill, Sarah C. Botwinick, Jonathan L. Gonzales, Kevin Ho, Allison Kundrik, Luke Kwon, Tim LaRowe, Khanh P. Nguyen, Fredy Ramirez, Peter Schaedler, William Ulrich, Amber Wallace, Yuchen Wan, and Benjamin Weinfeld. 2018. What Would You Do? Design Fiction and Ethics. In *Proceedings of the 2018 ACM International Conference on Supporting Group Work* (Sanibel Island, Florida, USA) (GROUP '18). Association for Computing Machinery, New York, NY, USA, 244–256. <https://doi.org/10.1145/3148330.3149405>
- [13] Katharina-Maria Behr, Andreas Nosper, Christoph Klimmt, and Tilo Hartmann. 2005. Some Practical Considerations of Ethical Issues in VR Research. *Presence: Teleoper. Virtual Environ.* 14, 6 (dec 2005), 668–676. <https://doi.org/10.1162/105474605775196535>
- [14] Michael S. Bernstein, Margaret Levi, David Magnus, Betsy A. Rajala, Debra Satz, and Quinn Waeiss. 2021. Ethics and society review: Ethics reflection as a precondition to research funding. *Proceedings of the National Academy of Sciences* 118, 52 (2021), e2117261118. <https://doi.org/10.1073/pnas.2117261118> arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.2117261118>
- [15] Alina Beygelzimer, Yann Dauphin, Percy Liang, and Jennifer Wortman Vaughan. 2021. *Introducing the NeurIPS 2021 Paper Checklist*. NeurIPS. <https://neuripsconf.medium.com/introducing-the-neurips-2021-paper-checklist-3220d6df500b>
- [16] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Iason Gabriel, and Shakir Mohamed. 2022. Power to the People? Opportunities and Challenges for Participatory AI. In *Equity and Access in Algorithms, Mechanisms, and Optimization* (Arlington, VA, USA) (EAAMO '22). Association for Computing Machinery, New York, NY, USA, Article 6, 8 pages. <https://doi.org/10.1145/3551624.3555290>
- [17] Pernille Bjorn, Casey Fiesler, Michael Muller, Jessica Pater, and Pamela Wisniewski. 2018. Research Ethics Town Hall Meeting. In *Proceedings of the 2018 ACM International Conference on Supporting Group Work* (Sanibel Island, Florida, USA) (GROUP '18). Association for Computing Machinery, New York, NY, USA, 393–396. <https://doi.org/10.1145/3148330.3154523>
- [18] Julian Blecker. 2022. *Design Fiction*. John Wiley I& Sons, Ltd, Berlin, Chapter 24, 561–578. <https://doi.org/10.1002/9781119815075.ch47> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781119815075.ch47>
- [19] Dounia Bourabain. 2021. Everyday sexism and racism in the ivory tower: The experiences of early career researchers on the intersection of gender and ethnicity in the academic workplace. *Gender, Work & Organization* 28, 1 (2021), 248–267.
- [20] Stacy M. Branham, Anja Thieme, Lisa P. Nathan, Steve Harrison, Deborah Tatar, and Patrick Olivier. 2014. Co-Creating I&I: Identity-Making in CSCW: Revisiting Ethics in Design Research. In *Proceedings of the Companion Publication of the 17th ACM Conference on Computer Supported Cooperative Work I& Social Computing* (Baltimore, Maryland, USA) (CSCW Companion '14). Association for Computing Machinery, New York, NY, USA, 305–308. <https://doi.org/10.1145/2556420.2558859>
- [21] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [22] Barry Brown, Alexandra Weilenmann, Donald McMillan, and Airi Lampinen. 2016. Five Provocations for Ethical HCI Research. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 852–863. <https://doi.org/10.1145/2858036.2858313>
- [23] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., NeurIPS 2020, 1877–1901. <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>
- [24] Amy S. Bruckman, Casey Fiesler, Jeff Hancock, and Cosmin Munteanu. 2017. CSCW Research Ethics Town Hall: Working Towards Community Norms. In

- Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland, Oregon, USA) (CSCW '17 Companion). Association for Computing Machinery, New York, NY, USA, 113–115. <https://doi.org/10.1145/3022198.3022199>
- [25] Federico Cabitza, Raffaele Rasoini, and Gian Franco Gensini. 2017. Unintended consequences of machine learning in medicine. *Jama* 318, 6 (2017), 517–518.
- [26] Shruthi Sai Chivukula, Ziqing Li, Anne C. Pivonka, Jingning Chen, and Colin M. Gray. 2021. Surveying the Landscape of Ethics-Focused Design Methods. *CoRR* abs/2102.08909 (2021), 1–23. [arXiv:2102.08909](https://arxiv.org/abs/2102.08909) <https://arxiv.org/abs/2102.08909>
- [27] Andy Cockburn, Carl Gutwin, and Alan Dix. 2018. HARK No More: On the Preregistration of CHI Experiments. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3173574.3173715>
- [28] Natasha Crampton. 2022. Microsoft's framework for building AI systems responsibly. <https://blogs.microsoft.com/on-the-issues/2022/06/21/microsofts-framework-for-building-ai-systems-responsibly/>
- [29] CVPR. 2022. CVPR Ethics Guidelines. <https://cvpr2022.thecvf.com/ethics-guidelines>
- [30] Michela Del Vicario, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. 2016. The spreading of misinformation online. *Proceedings of the National Academy of Sciences* 113, 3 (2016), 554–559.
- [31] Anthony Dunne and Fiona Raby. 2013. *Speculative Everything: Design, Fiction, and Social Dreaming*. The MIT Press, NY, USA. <http://www.jstor.org/stable/j.ctt9qf7j7>
- [32] Elsie L. Echeverri-Carroll, Michael D Oden, David V Gibson, and Evan A Johnston. 2018. Unintended consequences on gender diversity of high-tech growth and labor market polarization. *Research Policy* 47, 1 (2018), 209–217.
- [33] Rose Eveleth. 2018. Google Glass wasn't a failure, it raised crucial concerns. <https://www.wired.com/story/google-glass-reasonable-expectation-of-privacy/>
- [34] Casey Fiesler. 2019. Ethical considerations for research involving (speculative) public data. *Proceedings of the ACM on Human-Computer Interaction* 3, GROUP (2019), 1–13.
- [35] Casey Fiesler, Melissa Densmore, Michael Muller, and Cosmin Munteanu. 2021. SIGCHI Research Ethics Committee Town Hall. In *Companion Publication of the 2021 Conference on Computer Supported Cooperative Work and Social Computing* (Virtual Event, USA) (CSCW '21). Association for Computing Machinery, New York, NY, USA, 232–233. <https://doi.org/10.1145/3462204.3483283>
- [36] Casey Fiesler, Natalie Garrett, and Nathan Beard. 2020. What Do We Teach When We Teach Tech Ethics? A Syllabi Analysis. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education* (Portland, OR, USA) (SIGCSE '20). Association for Computing Machinery, New York, NY, USA, 289–295. <https://doi.org/10.1145/3328778.3366825>
- [37] Casey Fiesler, Alyson Young, Tamara Peyton, Amy S. Bruckman, Mary Gray, Jeff Hancock, and Wayne Lutters. 2015. Ethics for Studying Online Sociotechnical Systems in a Big Data World. In *Proceedings of the 18th ACM Conference Companion on Computer Supported Cooperative Work & Social Computing* (Vancouver, BC, Canada) (CSCW'15 Companion). Association for Computing Machinery, New York, NY, USA, 289–292. <https://doi.org/10.1145/2685553.2685558>
- [38] Office for Human Research Protections (OHRP). 2022. 45 CFR 46. Office for Human Research Protections (OHRP). <https://www.hhs.gov/ohrp/regulations-and-policy/regulations/45-cfr-46/index.html>
- [39] Christopher Frauenberger, Amy S. Bruckman, Cosmin Munteanu, Melissa Densmore, and Jenny Waycott. 2017. Research Ethics in HCI: A Town Hall Meeting. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI EA '17). Association for Computing Machinery, New York, NY, USA, 1295–1299. <https://doi.org/10.1145/3027063.3051135>
- [40] Batya Friedman and David G. Hendry. 2019. *Value Sensitive Design: Shaping Technology with Moral Imagination*. The MIT Press, Cambridge, MA, USA. <https://doi.org/10.7551/mitpress/7585.001.0001>
- [41] Batya Friedman, Peter H Kahn, and Alan Borning. 2008. Value sensitive design and information systems. *The handbook of information and computer ethics* 1, 1 (2008), 69–101. <https://doi.org/10.1002/9780470281819.ch4>
- [42] Batya Friedman, Lisa Nathan, Shaun Kane, and John Lin. 2018. Envisioning Cards. <https://www.envisioningcards.com/>
- [43] Ian G Johnson and Clara Crivellaro. 2021. Opening Research Commissioning To Civic Participation: Creating A Community Panel To Review The Social Impact of HCI Research Proposals. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 597, 17 pages. <https://doi.org/10.1145/3411764.3445113>
- [44] Ajit G. Pillai, A. Baki Kocaballi, Tuck Wah Leong, Rafael A. Calvo, Nassim Parvin, Katie Shilton, Jenny Waycott, Casey Fiesler, John C. Havens, and Naseem Ahmadpour. 2021. Co-Designing Resources for Ethics Education in HCI. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI EA '21). Association for Computing Machinery, New York, NY, USA, Article 109, 5 pages. <https://doi.org/10.1145/3411763.3441349>
- [45] Timmit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna M. Wallach, Hal Daumé III, and Kate Crawford. 2018. Datasheets for Datasets. *CoRR* abs/1803.09010 (2018), 1–11. [arXiv:1803.09010](https://arxiv.org/abs/1803.09010) <https://arxiv.org/abs/1803.09010>
- [46] Colin M. Gray and Shruthi Sai Chivukula. 2019. Ethical Mediation in UX Practice. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3290605.3300408>
- [47] Travis Greene, Amit Dhurandhar, and Galit Shmueli. 2022. Atomist or Holist? A Diagnosis and Vision for More Productive Interdisciplinary AI Ethics Dialogue. <https://doi.org/10.48550/ARXIV.2208.09174>
- [48] Barbara Grimpe, Mark Hartswood, and Marina Jirotko. 2014. Towards a Closer Dialogue between Policy and Practice: Responsible Design in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) (CHI '14). Association for Computing Machinery, New York, NY, USA, 2965–2974. <https://doi.org/10.1145/2556288.2557364>
- [49] Oliver L. Haimson and Anna Lauren Hoffmann. 2016. Constructing and enforcing "authentic" identity online: Facebook, real names, and non-normative identities. *First Monday* 21, 6 (Jun. 2016), 1–6. <https://doi.org/10.5210/fm.v21i6.6791>
- [50] Sara Hajian, Francesco Bonchi, and Carlos Castillo. 2016. Algorithmic Bias: From Discrimination Discovery to Fairness-Aware Data Mining. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Francisco, California, USA) (KDD '16). Association for Computing Machinery, New York, NY, USA, 2125–2126. <https://doi.org/10.1145/2939672.2945386>
- [51] Eric Hand. 2010. Citizen science: People power. *Nature News* 466, 7307 (2010), 685–687.
- [52] David Hankerson, Andrea R. Marshall, Jennifer Booker, Houda Elmimouni, Imani Walker, and Jennifer A. Rode. 2016. Does Technology Have Race?. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (San Jose, California, USA) (CHI EA '16). Association for Computing Machinery, New York, NY, USA, 473–486. <https://doi.org/10.1145/2851581.2892578>
- [53] Christina N. Harrington, Shamika Klassen, and Yolanda A. Rankin. 2022. "All That You Touch, You Change": Expanding the Canon of Speculative Design Towards Black Futuring. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 450, 10 pages. <https://doi.org/10.1145/3491102.3502118>
- [54] Michael I Harrison, Ross Koppel, and Shirly Bar-Lev. 2007. Unintended consequences of information technologies in health care—an interactive sociotechnical analysis. *Journal of the American medical informatics Association* 14, 5 (2007), 542–549.
- [55] Brent J. Hecht, Lauren Wilcox, Jeffrey P. Bigham, Johannes Schöning, Ehsan Hoque, Jason Ernst, Yonatan Bisk, Luigi De Russis, Lana Yarosh, Bushra Anjum, Danish Contractor, and Cathy Wu. 2021. It's Time to Do Something: Mitigating the Negative Impacts of Computing Through a Change to the Peer Review Process. *CoRR* abs/2112.09544 (2021), 1–6. [arXiv:2112.09544](https://arxiv.org/abs/2112.09544) <https://arxiv.org/abs/2112.09544>
- [56] Dirk Hovy, Shannon Spruit, Margaret Mitchell, Emily M. Bender, Michael Strube, and Hanna Wallach (Eds.). 2017. *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*. Association for Computational Linguistics, Valencia, Spain. <https://doi.org/10.18653/v1/W17-16>
- [57] Samuel P Huntington. 1971. The change to change: Modernization, development, and politics. *Comparative politics* 3, 3 (1971), 283–322.
- [58] ICSME 2021. 2021. ICSME Registered Reports Track. <https://icsme2021.github.io/cfp/RegisteredReportsTrack.html>, last accessed: September 6, 2021.
- [59] Brandon Ingram, Daniel Jones, Andrew Lewis, Matthew Richards, Charles Rich, and Lance Schachterle. 2010. A Code of Ethics for Robotics Engineers. In *Proceedings of the 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI '10)*. IEEE Press, Osaka, Japan, 103–104.
- [60] Lilly Irani, Janet Vertesi, Paul Dourish, Kavita Philip, and Rebecca E. Grinter. 2010. Postcolonial Computing: A Lens on Design and Development. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '10). Association for Computing Machinery, New York, NY, USA, 1311–1320. <https://doi.org/10.1145/1753326.1753522>
- [61] Liwei Jiang. 2021. Towards machine ethics and norms. <https://blog.allenai.org/towards-machine-ethics-and-norms-d64f2bdde6a3>
- [62] Deborah G Johnson. 1985. Computer ethics. *Englewood Cliffs (NJ)* 10 (1985), 102926.
- [63] Katarina Kertysova. 2018. Artificial Intelligence and Disinformation: How AI Changes the Way Disinformation is Produced, Disseminated, and Can Be Countered. *Security and Human Rights* 29, 1-4 (2018), 55–81.
- [64] Hans K. Klein and Daniel Lee Kleinman. 2002. The Social Construction of Technology: Structural Considerations. *Science, Technology, & Human Values* 27

- (2002), 28–52.
- [65] Rob Kling. 1991. Computerization and social transformations. *Science, Technology, & Human Values* 16, 3 (1991), 342–367.
- [66] Rob Kling. 1996. *Computerization and controversy: Value conflicts and social choices*. Elsevier, Amsterdam, Netherlands.
- [67] Frank Hyneman Knight. 1921. *Risk, uncertainty and profit*. Vol. 31. Houghton Mifflin, Wilmington, Delaware, USA.
- [68] Allison Koenecke, Andrew Nam, Emily Lake, Joe Nudell, Minnie Quartey, Zion Mengesha, Connor Toups, John R Rickford, Dan Jurafsky, and Sharad Goel. 2020. Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences* 117, 14 (2020), 7684–7689.
- [69] Kushwanth Koya and Gobinda Chowdhury. 2020. Measuring Impact of Academic Research in Computer and Information Science on Society. In *Proceedings of the 2020 2nd Asia Pacific Information Technology Conference* (Bali Island, Indonesia) (APIT 2020). Association for Computing Machinery, New York, NY, USA, 78–85. <https://doi.org/10.1145/3379310.3379312>
- [70] Robert Kraut, Michael Patterson, Vicki Lundmark, Sara Kiesler, Tridas Mukophadhayay, and William Scherlis. 1998. Internet paradox: A social technology that reduces social involvement and psychological well-being? *American psychologist* 53, 9 (1998), 1017.
- [71] Andrey Kurenkov. 2022. Lessons from the gpt-4chan controversy. <https://thegradient.pub/gpt-4chan-lessons/>
- [72] Peter Lee. 2016. Learning from Tay's introduction. <https://blogs.microsoft.com/blog/2016/03/25/learning-tays-introduction/>
- [73] Joseph Lindley and Paul Coulton. 2016. 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* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 4032–4043. <https://doi.org/10.1145/2858036.2858446>
- [74] Joseph Lindley, Paul Coulton, and Miriam Sturdee. 2017. Implications for Adoption. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 265–277. <https://doi.org/10.1145/3025453.3025742>
- [75] Sebastian Linxen, Christian Sturm, Florian Brühlmann, Vincent Cassau, Klaus Opwis, and Katharina Reinecke. 2021. How WEIRD is CHI?. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 143, 14 pages. <https://doi.org/10.1145/3411764.3445488>
- [76] David Liu, Priyanka Nanayakkara, Sarah Ariyan Sakha, Grace Abuhamad, Su Lin Blodgett, Nicholas Diakopoulos, Jessica R. Hullman, and Tina Eliassi-Rad. 2022. Examining Responsibility and Deliberation in AI Impact Statements and Ethics Reviews. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* (Oxford, United Kingdom) (AI/ES '22). Association for Computing Machinery, New York, NY, USA, 424–435. <https://doi.org/10.1145/3514094.3534155>
- [77] Octavian Mihai Machidon. 2018. Societal implications of current and emerging "smart" technologies. In *Smart Technologies: Breakthroughs in Research and Practice*. IGI Global, Country unknown/Code not available, 305–317.
- [78] Wendy E. Mackay. 1995. Ethics, Lies and Videotape... In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '95). ACM Press/Addison-Wesley Publishing Co., USA, 138–145. <https://doi.org/10.1145/223904.223922>
- [79] Michael A. Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. 2020. Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376445>
- [80] Donald McMillan, Alistair Morrison, and Matthew Chalmers. 2013. Categorized Ethical Guidelines for Large Scale Mobile HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) (CHI '13). Association for Computing Machinery, New York, NY, USA, 1853–1862. <https://doi.org/10.1145/2470654.2466245>
- [81] Amanda Menking, Ingrid Erickson, and Wanda Pratt. 2019. People Who Can Take It: How Women Wikipedians Negotiate and Navigate Safety. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). ACM, New York, NY, USA, Article 472, 14 pages. <https://doi.org/10.1145/3290605.3300702>
- [82] Robert K Merton. 1936. The unanticipated consequences of purposive social action. *American sociological review* 1, 6 (1936), 894–904.
- [83] Jacob Metcalf, Emanuel Moss, et al. 2019. Owning Ethics: Corporate Logics, Silicon Valley, and the Institutionalization of Ethics. *Social Research: An International Quarterly* 86, 2 (2019), 449–476.
- [84] Rada Mihalcea. 2021. ACL establishes its Ethics Committee. <https://www.aclweb.org/portal/content/acl-establishes-its-ethics-committee>
- [85] Josh Aaron Miller, Firas Khatib, Haley Hammond, Seth Cooper, and Scott Horowitz. 2020. Introducing Foldit education mode. *Nature Structural & Molecular Biology* 27, 9 (2020), 769–770.
- [86] James H Moor. 1985. What is computer ethics? *Metaphilosophy* 16, 4 (1985), 266–275.
- [87] Carol Moser, Sarita Y. Schoenebeck, and Katharina Reinecke. 2016. Technology at the Table: Attitudes About Mobile Phone Use at Mealtimes. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI) (Santa Clara, California, USA) (CHI '16). ACM, New York, NY, USA, 1881–1892. <https://doi.org/10.1145/2858036.2858357>
- [88] Cosmin Munteanu, Amy Bruckman, Michael Muller, Christopher Frauenberger, Casey Fiesler, Robert E. Kraut, Katie Shilton, and Jenny Waycott. 2019. SIGCHI Research Ethics Town Hall. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI EA '19). Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3290607.3311742>
- [89] Priyanka Nanayakkara, Nicholas A. Diakopoulos, and Jessica R. Hullman. 2020. Anticipatory Ethics and the Role of Uncertainty. *ArXiv abs/2011.13170* (2020), 1–4.
- [90] Priyanka Nanayakkara, Jessica Hullman, and Nicholas Diakopoulos. 2021. Unpacking the Expressed Consequences of AI Research in Broader Impact Statements. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (Virtual Event, USA) (AI/ES '21). Association for Computing Machinery, New York, NY, USA, 795–806. <https://doi.org/10.1145/3461702.3462608>
- [91] 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* (Cape Town, South Africa) (DIS '08). Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/1394445.1394446>
- [92] Openai. 2022. Dalle-2-preview/system-card.md at main · openai/dalle-2-preview. <https://github.com/openai/dalle-2-preview/blob/main/system-card.md>
- [93] Yuren Pang, Katharina Reinecke, and René Just. 2022. Apéritif: Scaffolding Preregistrations to Automatically Generate Analysis Code and Methods Descriptions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 207, 15 pages. <https://doi.org/10.1145/3491102.3517707>
- [94] Donn B. Parker, Susan Swope, Bruce N. Baker, and Eric A. Weiss. 1995. *All in a Day's Work: Nine Provocative Examples in the Practice of Computing Professionals*. Academic Press, Inc., USA, 870–875.
- [95] Nassim Parvin and Anne Pollock. 2020. Unintended by Design: On the Political Uses of "Unintended Consequences". *Engaging Science, Technology, and Society* 6 (2020), 320–327.
- [96] Robert A. Pearlman. 2021. *Ethics committees and consultation*. University of Washington. Retrieved 2021-09-14 from <https://depts.washington.edu/bhdept/ethics-medicine/bioethics-topics/detail/64>
- [97] Karime Pereira and M Greff. 2020. Diversity in robotics: From diverse teams to diverse impact.
- [98] Carina EA Prunkl, Carolyn Ashurst, Markus Anderljung, Helena Webb, Jan Leike, and Allan Dafoe. 2021. Institutionalizing ethics in AI through broader impact requirements. *Nature Machine Intelligence* 3, 2 (2021), 104–110.
- [99] Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT* '20). Association for Computing Machinery, New York, NY, USA, 33–44. <https://doi.org/10.1145/3351095.3372873>
- [100] Erick Jose Ramirez and Scott Labarge. 2018. Real Moral Problems in the Use of Virtual Reality. *Ethics and Inf. Technol.* 20, 4 (dec 2018), 249–263. <https://doi.org/10.1007/s10676-018-9473-5>
- [101] Katharina Reinecke and Krzysztof Z. Gajos. 2015. LabintheWild: Conducting Large-Scale Online Experiments With Uncompensated Samples. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (Vancouver, BC, Canada) (CSCW '15). Association for Computing Machinery, New York, NY, USA, 1364–1378. <https://doi.org/10.1145/2675133.2675246>
- [102] Bradford W Reynolds, Melissa W Burek, Billy Henson, and Bonnie S Fisher. 2013. The unintended consequences of digital technology: Exploring the relationship between sexting and cybervictimization. *Journal of Crime and Justice* 36, 1 (2013), 1–17.
- [103] Ronald E. Robertson, Shan Jiang, Kenneth Joseph, Lisa Friedland, David Lazer, and Christo Wilson. 2018. Auditing Partisan Audience Bias within Google Search. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 148 (nov 2018), 22 pages. <https://doi.org/10.1145/3274417>
- [104] Everett M. Rogers. 1976. New Product Adoption and Diffusion. *Journal of Consumer Research* 2 (1976), 290–301.
- [105] Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. The Risk of Racial Bias in Hate Speech Detection. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for

- Computational Linguistics, Florence, Italy, 1668–1678. <https://doi.org/10.18653/v1/P19-1163>
- [106] Hong Shen, Wesley H. Deng, Aditi Chattopadhyay, Zhiwei Steven Wu, Xu Wang, and Haiyi Zhu. 2021. Value Cards: An Educational Toolkit for Teaching Social Impacts of Machine Learning through Deliberation. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (*FAccT '21*). Association for Computing Machinery, New York, NY, USA, 850–861. <https://doi.org/10.1145/3442188.3445971>
- [107] Katie Shilton. 2018. Values and Ethics in Human-Computer Interaction. *Found. Trends Hum. Comput. Interact.* 12 (2018), 107–171.
- [108] Pao Siangliulue, Kenneth C. Arnold, Krzysztof Z. Gajos, and Steven P. Dow. 2015. Toward Collaborative Ideation at Scale: Leveraging Ideas from Others to Generate More Creative and Diverse Ideas. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work I& Social Computing* (Vancouver, BC, Canada) (*CSCW '15*). Association for Computing Machinery, New York, NY, USA, 937–945. <https://doi.org/10.1145/2675133.2675239>
- [109] Robert Soden, Michael Skirpan, Casey Fiesler, Zahra Ashktorab, Eric P. S. Baumer, Mark Blythe, and Jasmine Jones. 2019. CHI4EVL: Creative Speculation on the Negative Impacts of HCI Research. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (*CHI EA '19*). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3290607.3299033>
- [110] Kate Starbird. 2017. Examining the Alternative Media Ecosystem Through the Production of Alternative Narratives of Mass Shooting Events on Twitter. *Proceedings of the International AAAI Conference on Web and Social Media* 11, 1 (May 2017), 230–239. <https://doi.org/10.1609/icwsm.v11i1.14878>
- [111] Kate Starbird. 2019. Disinformation's spread: bots, trolls and all of us. *Nature* 571, 449 (2019), 449. <https://par.nsf.gov/biblio/10170694>
- [112] Amanda Stent. 2022. *Guidelines for ethics reviewing*. Association for Computational Linguistics. <https://aclrollingreview.org/ethicsreviewtutorial>
- [113] Constantine Stephanidis, Gavriel Salvendy, Margherita Antona, Jessie YC Chen, Jianming Dong, Vincent G Duffy, Xiaowen Fang, Cali Fidopiastis, Gino Fragoneni, Limin Paul Fu, et al. 2019. Seven HCI grand challenges. *International Journal of Human-Computer Interaction* 35, 14 (2019), 1229–1269.
- [114] Julia Stoyanovich, Serge Abiteboul, Bill Howe, HV Jagadish, and Sebastian Schelter. 2022. Responsible data management. *Commun. ACM* 65, 6 (2022), 64–74.
- [115] Miriam Sturdee, Joseph Lindley, Conor Linehan, Chris Elsdon, Neha Kumar, Tawanna Dillahunt, Regan Mandryk, and John Vines. 2021. Consequences, Schmonsequences! Considering the Future as Part of Publication and Peer Review in Computing Research. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (*CHI EA '21*). Association for Computing Machinery, New York, NY, USA, Article 95, 4 pages. <https://doi.org/10.1145/3411763.3441330>
- [116] Karl-Erik Sveiby, Pernilla Gripenberg, Beata Segercrantz, Andreas Eriksson, and Alexander Aminoff. 2009. Unintended and undesirable consequences of innovation. In *Unintended and undesirable consequences of innovation*. XX ISPIIM conference, The Future of Innovation. Vienna, Country unknown/Code not available, 1.
- [117] Naoya Tojo, Haruka Yoshida, Tomoko Oto, and Sumaru Niida. 2022. How Can We Expand Unintended Outcomes in Participatory Design?. In *Proceedings of the Participatory Design Conference 2022 - Volume 2* (Newcastle upon Tyne, United Kingdom) (*PDC '22*). Association for Computing Machinery, New York, NY, USA, 191–195. <https://doi.org/10.1145/3537797.3537873>
- [118] Kentaro Toyama. 2015. *Geek heresy: Rescuing social change from the cult of Technology*. Perseus Books Group, New York, NY.
- [119] Khai Truong. 2017. Pilot Studies: When and How to Conduct Them When Conducting User Studies. *GetMobile: Mobile Comp. and Comm.* 20, 4 (apr 2017), 8–11. <https://doi.org/10.1145/3081016.3081020>
- [120] Cora van Leeuwen, Shirley A. Elprama, An Jacobs, Rob Heyman, Jo Pierson, and Pieter Duysburgh. 2020. Unethically Me: Explaining Artificial Intelligence's Results by Being Unethical. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society* (Tallinn, Estonia) (*NordiCHI '20*). Association for Computing Machinery, New York, NY, USA, Article 139, 3 pages. <https://doi.org/10.1145/3419249.3420065>
- [121] Dieter Vanderelst and Alan Winfield. 2018. The Dark Side of Ethical Robots. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society* (New Orleans, LA, USA) (*AIES '18*). Association for Computing Machinery, New York, NY, USA, 317–322. <https://doi.org/10.1145/3278721.3278726>
- [122] Lucas Nunes Vieira, Minako O'Hagan, and Carol O'Sullivan. 2021. Understanding the societal impacts of machine translation: a critical review of the literature on medical and legal use cases. *Information, Communication & Society* 24, 11 (2021), 1515–1532.
- [123] Jessica Vitak, Nicholas Proferes, Katie Shilton, and Zahra Ashktorab. 2017. Ethics regulation in social computing research: Examining the role of institutional review boards. *Journal of Empirical Research on Human Research Ethics* 12, 5 (2017), 372–382.
- [124] Jessica Vitak, Katie Shilton, and Zahra Ashktorab. 2016. Beyond the Belmont Principles: Ethical Challenges, Practices, and Beliefs in the Online Data Research Community. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work I& Social Computing* (San Francisco, California, USA) (*CSCW '16*). Association for Computing Machinery, New York, NY, USA, 941–953. <https://doi.org/10.1145/2818048.2820078>
- [125] Greg Walsh. 2018. Towards Equity and Equality in American Co-Design: A Case Study. In *Proceedings of the 17th ACM Conference on Interaction Design and Children* (Trondheim, Norway) (*IDC '18*). Association for Computing Machinery, New York, NY, USA, 434–440. <https://doi.org/10.1145/3202185.3202768>
- [126] Yilun Wang and Michal Kosinski. 2018. Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of personality and social psychology* 114, 2 (2018), 246.
- [127] Jenny Waycott, Cosmin Munteanu, Hilary Davis, Anja Thieme, Wendy Moncur, Roisin McNaney, John Vines, and Stacy Branham. 2016. Ethical Encounters in Human-Computer Interaction. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (San Jose, California, USA) (*CHI EA '16*). Association for Computing Machinery, New York, NY, USA, 3387–3394. <https://doi.org/10.1145/2851581.2856498>
- [128] David Gray Widder, Dawn Nafus, Laura Dabbish, and James Herbsleb. 2022. Limits and Possibilities for “Ethical AI” in Open Source: A Study of Deepfakes. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (*FAccT '22*). Association for Computing Machinery, New York, NY, USA, 2035–2046. <https://doi.org/10.1145/3531146.3533779>
- [129] Max L. Wilson, Wendy Mackay, Ed Chi, Michael Bernstein, Dan Russell, and Harold Thimbleby. 2011. RepliCHI - CHI Should Be Replicating and Validating Results More: Discuss. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems* (Vancouver, BC, Canada) (*CHI EA '11*). Association for Computing Machinery, New York, NY, USA, 463–466. <https://doi.org/10.1145/1979742.1979491>
- [130] Langdon Winner. 1980. Do Artifacts Have Politics? *Daedalus* 109, 1 (1980), 121–136. <http://www.jstor.org/stable/20024652>
- [131] Langdon Winner. 1986. *The Whale and the Reactor: A Search for Limits in an Age of High Technology*. University of Chicago Press, Chicago, IL.
- [132] Richmond Y. Wong. 2021. Using Design Fiction Memos to Analyze UX Professionals' Values Work Practices: A Case Study Bridging Ethnographic and Design Futuring Methods. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (*CHI '21*). Association for Computing Machinery, New York, NY, USA, Article 93, 18 pages. <https://doi.org/10.1145/3411764.3445709>
- [133] Richmond Y Wong, Deirdre K Mulligan, Ellen Van Wyk, James Pierce, and John Chuang. 2017. Eliciting values reflections by engaging privacy futures using design workbooks. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 1–26.
- [134] Daisy Yoo. 2017. Stakeholder Tokens: A Constructive Method for Value Sensitive Design Stakeholder Analysis. In *Proceedings of the 2017 ACM Conference Companion Publication on Designing Interactive Systems* (Edinburgh, United Kingdom) (*DIS '17 Companion*). Association for Computing Machinery, New York, NY, USA, 280–284. <https://doi.org/10.1145/3064857.3079161>