

Ethical Considerations for Research Involving (Speculative) Public Data

CASEY FIESLER, University of Colorado Boulder, USA

As the process of creating and sharing data about ourselves becomes more prevalent, researchers have access to increasingly rich data about human behavior. Framed as a fictional paper published at some point in the not-so-distant future, this design fiction draws from current inquiry and debate into the ethics of using public data for research, and speculatively extends this conversation into even more robust and more personal data that could exist when we design new technologies in the future. By looking to how the precedents of today might impact the practices of tomorrow, we can consider how we might design policies, ethical guidelines, and technologies that are forward-thinking.

CCS Concepts: • **Security and privacy** → *Social aspects of security and privacy*;

Additional Key Words and Phrases: design fiction; ethics; lifelogging; privacy; public data; social computing; research ethics; research methods; quantified self

ACM Reference Format:

Casey Fiesler. 2019. Ethical Considerations for Research Involving (Speculative) Public Data. *Proc. ACM Hum.-Comput. Interact.* 3, GROUP, Article 249 (December 2019), 13 pages. <https://doi.org/10.1145/3370271>

1 Author's Introductory Notes

High profile cases in recent years such as the Facebook emotional contagion study [33], university researchers collecting facial recognition data from students without consent [31], and photos mined from dating sites used to train algorithms to predict sexual orientation [52] have challenged traditional notions of research ethics, particularly within the computing research community. Research ethics in terms of both regulation and training has focused largely on human subjects research, where harm is conceptualized via direct interaction with research participants [61]. However, the type of research now possible due to the scope and scale of publicly available data not only falls outside of these traditional regulatory schemes [71], but harms are often more indirect and less foreseeable. Within the human-computer interaction (HCI) research community, there have been calls to more carefully consider the possible negative consequences of our work [35], though this type of speculation about future harms can be challenging.

One possible avenue for addressing this challenge is through design fiction, which has been used to explore the potential outcomes of new design work, including the creation of fictional studies or papers that provide a space for critique [9]. Design fiction methods can encourage anticipation of and reflection about the potential downsides of technology design, research, and implementation [64]; indeed, reflection about speculative designs has been successful in encouraging technologists to consider how their practices might play out in the future [77]. Drawing from these ideas, the fictional paper that follows tackles the issue of research ethics and possible harms by exploring a

Author's address: Casey Fiesler, University of Colorado Boulder, Department of Information Science, Boulder, CO, 80309, USA, casey.fiesler@colorado.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-0142/2019/12-ART249 \$15.00

<https://doi.org/10.1145/3370271>

plausible technological future and the kinds of harms that might occur if today's ethical precedents expand to tomorrow. By scaffolding this type of creative speculation, the goal of this paper is to encourage researchers to think more broadly about the types of harms that are possible and the potential consequences of the ethical norms that are in the process of being created.

2 Introduction: New Ethical Considerations for Public Data

As the process of creating and sharing data about ourselves became more prevalent, researchers have had access to increasingly rich data about human behavior. Accordingly, ethical practices surrounding the use of such data has long been a topic of debate and inquiry even as the nature and scope of pervasive public data has changed due to evolving technology and practice. These conversations began with simple observations of people in public spaces, and when those public spaces moved online, the ethical dilemmas became even more complex. We have since moved beyond the precise concerns raised by early, simplistic instantiations of social sharing, such as once highly popular Internet platforms like Twitter (which involved largely text, with some basic static media) and YouTube (short, static, unconnected video content). However, many of the underlying principles remain, and the evolution of these ethical debates within research communities are still instructive for thinking through more contemporary challenges—specifically, regarding the availability of LifeStream data for research.

Though recent controversies related to LifeStream have been framed in the media as privacy issues, a subset of these actually have research ethics at their heart—notably, the labeling of a college student as depressed by a major newspaper, and the harassment of a group of women game developers via cruelly edited videos of deceased loved ones. These two cases touch on a number of pressing complexities when it comes to researchers using LifeStream data, including inference, amplification of content, false positives, use of deleted data, and sharing public datasets.

Through a historical lens, this paper interrogates these evolving ethical considerations for research that makes use of public LifeStream data. Using these case studies, and illustrating how past precedent has influenced current practices and discussions, we can better position current debates in order to inform ongoing ethical norms and guidelines as our research moves forward.

3 History and Related Work

3.1 Lifelogging

Lifelogging is a phenomenon where people passively record, typically through wearable devices, data from their daily lives in varying amounts of detail [32]. The concept has some overlap with quantified-self techniques, though rather than focusing on a single analytic such as health self-monitoring, lifelogging moves towards the totality of life experience [32]. Early large-scale examples of lifelogging included MyLifeBits and SemanticLife, both projects from the early 2000s that attempted to create digital archives of personal information [59]. Other early lifelogging-related technologies tended to focus on personal, specific functions—for example, tracking eating habits [69], creating records of infants' milestones [43], or supporting people with impaired memory [51]. The potential for digital memory was cast as almost utopic, with images of the potential for perfect recall, personal reflection, improved productivity, and safety [18].

In addition to concerns about information overload [18], privacy was frequently raised as a significant issue for the potential widespread adoption of lifelogging [14, 26, 37, 59]. Even without the public dissemination of lifelogging data that we associate with the practice today, privacy was seen as a roadblock due to legal issues and the recording of data beyond oneself (e.g., bystanders or sensitive settings [19, 26, 68]). The failure of Google Glass, the first significant commercial augmented reality product, was in part due to both concern over privacy law and widespread

public discomfort, illustrating the difficulty in navigating both policy and social norms [46, 74]. Automated methods for preserving privacy in-situ for lifelogging were developed quite early on—for example, defensive frameworks defined by physical location such as bedrooms and bathrooms [68], and methods based on facial and movement detection [69]. However, these privacy-related issues become even more pronounced when lifelogging data is shared beyond the user [26, 59].

Though lifelogging can focus on personal functions and does not always have a sharing component, social sharing is a natural extension of the phenomenon [26, 59], and indeed, some early studies of lifelogging for personal reflection revealed the desire to share more broadly [17, 43]. Though there are earlier examples of public-facing lifelogging (for example, through diaries or unsophisticated early wearables), the website JenniCAM, which began in 1996 and streamed a young woman’s life via webcam, is typically considered the first widely shared lifelog [41]. Though as lifelogging technologies evolved, sharing became a more integrated component, and early research regarding sharing preferences revealed an interest in sharing the majority of data, even amidst privacy concerns [37]. There has since been a great deal of research that has examined the privacy, ethical, and social implications of lifelogging generally [1, 14, 26], including implications for research ethics (particularly in the context of informed consent and scope of data [42, 48, 58]).

However, a significant difference between LifeStream and the lifelogging that took place ten years ago is the pervasive and public nature of LifeStream data, beyond even previously popular, static social media platforms such as Twitter and Instagram. Therefore, this paper focuses on one specific facet that is highly relevant for LifeStream: the implications for research ethics in the context of public, easily accessible lifelog data.

3.2 Public Data and Research Ethics

Discussion of the ethical implications of research involving public data notably surfaced when widespread use of the Internet as a research site introduced many complications to traditional notions of research ethics. In the United States, the Belmont Report (published in 1979) [66] has long been the major source of ethical standards for human subjects research, with universities largely governed by institutional review boards (IRBs). However, Internet research introduced some question into what even constitutes human subjects research. IRB definitions (still unchanged) require that the research involve “data through intervention or interaction with the individual” and/or “identifiable private information” [11]. Particularly with definitions of public spaces and public data on the Internet so unclear [10, 80], many IRBs have long considered the study of public data to not fall under their purview [11, 71]. However, even in the early days of these debates, researchers still acknowledged the potential harms that can occur with the study of online public data, particularly in the context of privacy violations [29, 56, 80].

Given the regulatory ambiguity and inconsistency [71], ongoing to this day, research communities have had to form their own ethical guidelines and norms for research involving public data. Similar to discussions of this issue with respect to social media data many years ago, researchers are struggling to create these norms, and there are differences in ethical heuristics from discipline to discipline and from researcher to researcher [72]. These differences are exacerbated by competing values—weighing individual and collective privacy concerns, for example, against innovation, progress on important problems, or replicability [44]. However, even as these debates continued, whether the data came from Facebook [80] or a dating site [81] or later LifeStream, the most commonly cited ethical heuristic for data collection remained whether the data is “public”—despite direct calls for this heuristic to change to include more nuance, such as the origin, characteristics, and purpose of the data collected [29].

Specific concerns raised over many years by researchers interrogating the ethical considerations of research using public data include the lack of informed consent [6, 40], identifiability of individuals

in data [5, 16, 80], the amplification of content beyond its intended context [22, 73], community-level harms [36], the implications of algorithmic inference [13], and considerations for sensitive contexts or populations [13, 22, 57], among others. However, these issues must be framed within the broader context of the benefits of scientific discovery; even when the public data available was simply static text on Twitter, researchers were able to make advancements that benefited, for example, public health [50, 62] and disaster response [47, 60]. For this reason, early explorations into framing the problem suggested that, in order to ensure that we do not shut down research with public data entirely, it is critical to create ethical guidelines that go beyond measures of publicness and consider the contextual factors at play that contribute to potential harms [29]. Though there have been a number of attempts since then to create such guidelines and codify them, and many individuals or even entire research communities have made progress in creating stable norms, there is still a great deal of uncertainty and disagreement with respect to ethical considerations for LifeStream data specifically.

4 LifeStream, Research, and Ethics

Over the past decades, a number of technological and social developments have rapidly altered the scope, scale, and character of data that researchers can now access. Whereas as early as the short-lived Google Glass project [74] and its precursors [37], people have had the ability to record their surroundings with little effort, the LifeStream camera was the first such technology to achieve truly significant uptake among the general public. Though many of the lifelogging technologies that preceded LifeStream focused on personal use, it was the first system designed specifically for social sharing. When first introduced, a number of commentators compared LifeStream to the technology introduced in Dave Eggers' 2013 novel *The Circle*, in which characters chose to “go transparent” and share every moment of their lives with the world, with first-person video constantly recorded, shared, and archived [23]—an idea that does in retrospect now seem quite prescient.

Though the precise number of active users is in contention, the userbase for LifeStream is comparable to that of Facebook at the height of its popularity. About the size of a dime, designed to stick easily to skin (typically worn just above the bridge of the nose), and requiring charging only once every 72 hours, the unobtrusive nature of the LifeStream camera was one factor that led to heavy initial adoption. Shifts in privacy norms in the years between the releases of Google Glass and LifeStream also made for easier public acceptance. In fact, normative changes continued to push adoption. A widespread values-shift towards accountability, fueled by a number of crises involving public figures in positions of power with society-level ramifications, heavily influenced increased adoption of LifeStream. As legal scholar Anita Allen predicted in 2003, the spread of social, economic, and political freedom, combined with increased ambivalence about privacy, has led to a culture in which individuals are expected to make public the more intimate details of their lives in a context in which they are also safer from punishment [1]. Interestingly, this profound cultural change also echoed themes introduced in *The Circle*, where social pressure contributed to politicians being among the first to go transparent. Some early enthusiastic LifeStream user communities even adopted mantras from the book as their own: Secrets Are Lies; Sharing Is Caring; Privacy Is Theft [23].

The LifeStream camera passively records first-person video constantly, except when deactivated manually or by the automated triggering of rules that a user can set—for example, based on detection of location, face, or heart rate. Research has shown that the most common rules include deactivation in bathrooms or bedrooms, when family members are in view, or when heart rate is low (as a proxy for sleep) or high (as a proxy for sexual activity). However, nearly 1/3 of all LifeStream users do not set any rules at all, which results in 24-hours-per-day archived video data, which is automatically uploaded to the sharing platform. Though shared data is public by default, users have the option

to use privacy settings that specify who can view their content (modeled after the “friends only” options from Facebook). However, similar to Facebook, users often have difficulty understanding or implementing privacy settings [27, 67], and users rarely use granular selective sharing or have different settings for different types of content [27]. Researchers have estimated that approximately 75% of all LifeStream data is publicly viewable by anyone.

Today’s widespread public sharing of LifeStream data without regard for highly salient privacy concerns might have come as a surprise to privacy researchers even a decade ago. However, the amount of data shared intentionally on social media at that time also might have shocked researchers from the decade before that, and generational differences in privacy attitudes—particularly regarding intentional social sharing—were already coming into view [34]. Moreover, there were many motivations for sharing personal information publicly even before the significant values shift to accountability over privacy—not just connection and encouragement [49] but even sharing private or stigmatized information in order to receive social support [2, 30] or identifiable information to seek help during a crisis [16].

This enormous trove of publicly available data has led to another “data gold rush” by researchers, reminiscent of the once dominance of Twitter data for research [25, 82]. This gold rush is also facilitated by an API that makes data collection an easy task without much computational work. As a result (again, as with Twitter [29]), LifeStream data is being used to answer innumerable different kinds of research questions across nearly every discipline imaginable. A quick search as of this writing reveals that within the last week, scholarly papers using LifeStream data have been published in journals for Linguistics (regional accent analysis), Gender Studies (gender roles in preschools as performed within and outside of a teacher’s presence), Media Studies (media consumption habits correlated to fanfiction production), Astronomy (stargazing data), Public Health (flu symptoms displayed by clinic walk-ins), and Human-Computer Interaction (behavioral analysis of LifeStream non-users). Finally, similar to the revelation from a meta-analysis of Twitter research in 2014 [82], none of these published papers make any mention of a university IRB or any other form of ethical review.

It is also worth noting that, though the focus here is on researchers’ uses of LifeStream data, there are a number of ethical concerns outside of this scope that also track to issues that once plagued social media—for example, the sharing and selling of user data with third parties without explicit consent beyond Terms of Service [28], manipulation of the LifeStream social feed for A/B testing without user knowledge [33], the internal collection of friends-only data to create user behavioral models [65], death and injury due to purposefully engaging in dangerous activity for LifeStream views [70], and opaque behavioral advertising techniques [24]. Additionally, beyond researchers, there are ethical concerns with other parties’ use of LifeStream data—notably, marketers and journalists. In addition to significant controversies (e.g., a data journalists’ analysis that revealed a number of Hollywood affairs among dyads of LifeStream non-users), there is also concern that again echoes a problem common with social media [20, 22]—the use of LifeStream data by journalists without knowledge or consent, to illustrate public opinion. Most recently, a news article about the historic cancellation of *The Simpsons* included of people reacting to the news—including one video that inadvertently outed a transgender woman.

In the realm of research ethics alone, there are a number of relevant and tricky considerations. Even a seemingly basic question like “is it okay that none of these research ‘participants’ are giving consent?” is highly complex. The next section will interrogate two specific case studies of real controversies, as a way of illustrating a subset of current challenges in considering the ethics of researching with LifeStream data.

4.1 Case Study 1: Inference, Amplification, and False Positives

At a computer science conference, a PhD student presented research that involved a novel LifeStream data mining method and algorithm for identifying college freshmen at risk of dropping out early. The presentation included a video montage of LifeStream data used in the project, and the video included a notable image of 19-year-old freshman Eliot looking into a mirror. The researcher highlighted this image as part of an analysis that flagged the user as having markers of depression, which was an important predictive feature in the algorithm. Though there were only perhaps 20 people attending the presentation, one of them was a journalist who was covering the conference. In a piece she wrote about this research, she included a few still images from the video to highlight the predictive features that the algorithm identified, including Eliot. Eliot's parents recognized his photo and were very alarmed; they pulled him out of school and sent him to a psychiatrist, insisting on in-patient treatment. According to the psychiatrist, he was not depressed—just homesick and lonely, but coping fairly well. Ironically, though he had been doing quite well in school prior to this situation, he then got behind due to his absence. The situation came to light because of the parents' own LifeStream feeds.

For the purposes of this discussion about researchers, set aside the actions of the journalist, though there are potential ethical issues with journalists highlighting public content no matter the source [20]. There has also been at least one case where a journalist not only quoted social media data from a research paper in a story, but even contacted the user [63]. However, a relevant question for the researcher is whether they should have considered the potential for the amplification of content beyond the 20 people in that room. This has long been a concern for the appearance of public data in research papers, where even when data was primarily text, it was common practice to provide verbatim quotes. In a 2018 study of articles using Twitter data that were published in PubMed, a database for health-related papers, 72% of those papers included at least one tweet, and of those, the user could be located 84% of the time [5]. In many cases, this may not be a problem—not only because so few people will read the paper, but also because the content might be unlikely to cause harm. For example, in the case of the tweet in the *New York Times* story, the tweet was about going fishing [63]. However, there were also research papers that used Twitter data to examine, for example, eating disorders [3] and HIV [79]. Perhaps public content on a sensitive topic like these (or depression) should be treated more carefully than content about fishing.

It is also worth noting that the content itself—the LifeStream video of Eliot looking into a mirror—did not involve Eliot saying, “I am depressed.” Instead, the label was placed on Eliot by an algorithm, which inferred his depression based on the sum of his LifeStream data. Even though these kinds of inferences tend to be more accurate than when they were based on tweets, they are still never 100% accurate. False positives that incorrectly identify the presence of a characteristic—particularly something private or even stigmatized like a mental health condition—can have significant negative consequences for the subject of the error [13], which makes the risk of amplification of the label even more severe. Though there are real benefits to this kind of research—the potential to correctly identify people who are at risk of harm and intervene in a way that can help them—there are risks as well, whether labels are accurate or inaccurate. In these cases, it is critical that researchers take steps to protect the subjects in their data.

This care is particularly important because LifeStream provides a complication beyond that of traditional social media: the presence of bystanders. LifeStream non-users still appear in this public data that is available for analysis, and even though they have not “signed up” for being included, they could be researched or labeled as well. For example, researchers have created algorithms that can identify someone's sexual orientation from their LifeStream feed very accurately, even if the data is sparse (unsurprising, considering that researchers once did the same with photographs

mined from a dating site [52]). It also is not unheard of for researchers to include photographs of people from public data in papers without their knowledge [38]. LifeStream data from a board game night is not equivalent to LifeStream data from a support group meeting, and it is imperative that researchers consider this context in order to make decisions about how they treat it.

4.2 Case Study 2: Deleted Data and Data Curation

A number of years ago, a 28-year-old North Carolina man committed suicide. Sometime last year, his widow, Mary, began receiving video messages: spliced together LifeStream content of her husband, edited to make it seem as if he were criticizing her and blaming her for his death—including footage of his suicide. Around the same time, three other women received similar messages, all cruelly remixed videos of deceased loved ones—a father, a sister, a child. All five were female game developers, and eventually it was revealed that the perpetrator was a self-described troll who had found a cruel new way to harass his victims. In Mary’s case, the video was even more shocking because after her husband’s death, she had requested that the LifeStream content related to his suicide be deleted, and the company had done so. The troll eventually revealed that he had obtained the video from a publicly available dataset compiled by a university researcher who had conducted a linguistic analysis of “last words”; the dataset included all LifeStream data from the last three days of over 5 million people who had died on the platform in the past 5 years, and it was published in a public repository in conjunction with the publication of a research paper.

A remaining significant tension with respect to research ethics is the tension between protecting data privacy and sharing data in the spirit of open science and replicability. Despite longstanding debates, there is still no clear resolution, particularly for data that is not public and may be identifiable [21], but there is a clear norm or even requirement within certain research communities (e.g., machine learning) to share datasets with published papers. There was a push to encourage this practice for public social media research, for a number of compelling reasons in favor of open science: to support reproducibility and comparability; to avoid digital divides in data accessibility; and to save time and money in the data collection process [75]. However, there are also ethical challenges that accompany this practice—for example, the difficulty of anonymization and possibility of re-identification of content creators, and the non-consensual archiving of user content that they might have deleted [75].

The “last words” dataset in this situation was, at the time of the researcher’s collection, all publicly available—all he did was curate and archive it. There are two issues here—first, the archiving of data that was later deleted, and second, the overall curation, regardless of whether the data was deleted. With respect to the first, a study of user attitudes about research ethics for Twitter data revealed that users largely considered the use of deleted tweets by researchers to be unethical [29], and LifeStream data is deleted more frequently than traditional social media was. The issue is one of respect for the user’s control over their own content; if they (or someone with power over their content, like a surviving relative in the case of a death) no longer want to share that content, should a researcher be able to take that control from them by creating an archive of it? Datasets are rarely cleaned for deleted or removed data, even for sensitive contexts [13]. However, one possibility for mitigating potential harm is to, rather than duplicating data, provide a dataset that includes pointers to where that data lives in the LifeStream feed; in this case, other researchers can recreate the dataset, absent any data that has been deleted [75].

The second question is whether there are ethical issues with simply the curation of the data, if all of it were public and available for anyone to find. A similar question came up in 2016, when a researcher released a public dataset of every profile from the dating site OKCupid without any anonymization [81]. The researcher’s response, when questioned about this, was that the data was “already public,” and that the dataset simply put it “in a more useful form” [81]. There was

considerable controversy following this response, particularly since, given the type of data that was available, it would have been trivial to connect that data (which included, for example, the answers to very personal questions about sexual preferences) to real names. Perhaps the sensitivity of the content matters; could the “last words” researcher have anticipated such a cruel use of the data he curated? Though sadly, using deceased loved ones as trolling ammunition is not unheard of [76], data regarding death is still more sensitive (and thus has more potential for harm) than many types of data. In a 2017 paper in which the data involved social media users who had committed suicide, the authors decided that because the dataset could “expose sensitive information widely,” rather than releasing the dataset publicly, they would make it available on a case-by-case basis, sharing only the subsets of data needed to fill a particular research request [39]. Particularly in cases where the context is potentially sensitive, this could be a good solution to balance harm mitigation and openness in research.

5 Conclusion

These two controversies highlight a number of the ethical challenges currently facing researchers who are making use of public LifeStream data. They also served to highlight this problem more widely—in the media, among researchers, in the general public—about researchers on LifeStream and what their ethical obligations should be. However, these discussions largely lack a historical perspective; this is nothing new. In fact, though these cases may seem extreme due to the personal nature and scope of the data involved, the complexities highlighted are very similar to those that researchers faced when collecting public social media data on platforms like Twitter and Facebook.

However, the practices of the past are still influencing those of the present. In much the same way that decisions in case law create legal precedent that can persist over decades of social and technological change, so too can ethical heuristics—like “but the data is already public.” The situations discussed here highlight how critical it is that researchers not only update their thinking to account for new technological and social arrangements but resist norms that are too rigid to account for context. Consider, for example, the legal doctrine of fair use, which provides exceptions to U.S. copyright law; this doctrine is designed as a set of flexibly weighed factors rather than bright-line rules in part so that it can adapt to changes in technical and content practices [4]. This has allowed the law, which was codified in the 1970s, to adapt to digital technology, the internet, and beyond.

However, though we have been asking similar questions around research ethics for many years, regulations are still ambiguous and inconsistently applied, practices still vary widely between disciplines and even individual researchers, and there is still not enough reflection about these issues among researchers to move us towards stable but contextually sensitive norms. Past precedent, codified by thousands of papers with the same methods, has helped to solidify our current position. Perhaps if we had worked more carefully and deliberately through the first ethical challenges that social media created for researchers, we would not still be asking the same questions at a point where the data we use is more robust and the consequences of our lack of forethought more severe.

6 Author’s Concluding Notes

The goal of this design fiction was to illustrate the importance of thinking through the ethical implications of research using public data now, even when the immediate harms may not be obvious. Speculative designs can encourage people to compare possible futures to the technologies and practices of the present [78]. Because the potential harms and privacy violations inherent in the availability of data like that envisioned in *The Circle* are more obvious than those of the present, my hope is that for researchers reading this paper, drawing a line between the present and a possible future illustrates the slippery slope between them. I drew inspiration from *The Circle* for LifeStream

for similar reasons as Wong et al. in their study that directly solicited values reflection based on technologies presented in the book [77, 78]; these technologies are already part of the public imagination (in part due to a film adaptation) and they represent a logical next step based on the trajectory of current technological development. Indeed, even as researchers today are creating design fiction to speculate about the technology of the future, those same designs are close to becoming reality [77].

Design fiction, as a way of considering possible futures, has long been perfectly positioned to interrogate the ethical implications of technology. It encourages us to explore consequences [8], whether we are considering radical digital interventions [53] or the natural social and ethical evolution of technology [55]. It also allows us to consider what it would mean for a technology to exist not just as a prototype but how it might be adopted in the real world [54]. Further, the design fiction presented here goes beyond how a technology might be adopted by users and considers instead how the data that it produces might be adopted and used by researchers. Design fiction can be seen as a kind of world building, where artefacts can serve as entry points into that world [15]; here, the entry points are research uses of data that illustrate particular types of consequences and harm, such as amplification and harassment.

Potential negative consequences of research and practice has recently been a topic of increased discussion within the HCI community, beyond just the design of new technology [35]. Just as we might ask “How does this work contribute to the future enslavement of humankind by evil robots?” [45], we should also ask what harm we might do with data when we forget that there are people behind it. Accordingly, this design fiction presents an example of a type of speculation that considers secondary rather than primary uses of a technology—which forces a focus on unintended consequences.

Unintended consequences are a critical part of ethical speculation, which is one reason that design fiction is well-suited for teaching ethics. Students are often not well-trained when it comes to research ethics—particularly researchers who are accustomed to working with data but not human subjects. It can be challenging to teach students to recognize larger ethical issues and the responsibility that comes with building or studying technology, but creative speculation is an excellent tool for engaging students in the scope and implications of current and future work in computer science [12]. It is hard to read a book like *The Circle* and not speculate about the ethical implications of this technology; if someone can think through the implications of the technology that someone might create in thirty years, they can also think through the implications of the technology that they might create (or study) tomorrow.

The implications explored in this design fiction were drawn from many discussions about open questions in research ethics for pervasive public data, largely from a string of workshops and panel discussions over the past five years. The types of harms explored, and the ethical decisions that might have helped mitigate them, track to suggestions in prior work for the types of context that researchers should be taking into account, particularly: the character of the content, the vulnerability of who created it, who will be making use of it, and what that use is [29]. A goal of design fiction is to allow visions to tomorrow to shape the research directions of today [7]. To add to this ongoing push for contextually driven research ethics decisions that move beyond “publicness” of data, my hope is that by looking to how the ethical precedents of today might impact the practices of tomorrow, we can work towards research, policies, ethical guidelines, and technologies that are forward-thinking.

Acknowledgments

This work was supported by NSF Award #1704303 as part of the PERVADE (Pervasive Data Ethics for Computational Research) collaborative research project.

References

- [1] Anita L. Allen. 2003. Privacy Isn't Everything: Accountability As a Personal and Social Good. *Alabama Law Review* 54 (2003), 1375–1391.
- [2] Nazanin Andalibi, Oliver L. Haimson, Munmun De Choudhury, and Andrea Forte. 2016. Understanding Social Media Disclosures of Sexual Abuse Through the Lenses of Support Seeking and Anonymity. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems* (2016), 3906–3918. <https://doi.org/10.1145/2858036.2858096>
- [3] Alina Arseniev-Koehler, Hedwig Lee, Tyler McCormick, and Megan A. Moreno. 2016. #Proana: Pro-Eating Disorder Socialization on Twitter. *Journal of Adolescent Health* 58, 6 (2016), 659–664. <https://doi.org/10.1016/j.jadohealth.2016.02.012>
- [4] Patricia Aufderheide and Peter Jaszi. 2011. *Reclaiming Fair Use: How to Put Balance Back Into Copyright*. University of Chicago Press, Chicago, IL.
- [5] John W. Ayers, Theodore L. Caputi, Camille Nebeker, and Mark Dredze. 2018. Don't quote me: reverse identification of research participants in social media studies. *npj Digital Medicine* 1, 1 (2018), 29–30. <https://doi.org/10.1038/s41746-018-0036-2>
- [6] S. Barocas and Helen Nissenbaum. 2014. Big data's end run around anonymity and consent. In *Privacy, big data, and the public good: Frameworks for engagement*, J. Lane, V. Stodden, S. Bender, and Helen Nissenbaum (Eds.). Cambridge University Press, Cambridge, UK, 44–75.
- [7] Eric P. S. Baumer, June Ahn, Mei Bie, Elizabeth M. Bonsignore, Ahmet Brutecene, Oguz Turan Buruk, Tamara Clegg, Allison Druin, Florian Echtler, Daniel Gruen, Mona Leigh Guha, Chelsea Hordatt, Antonio Kruger, Shachar Maidenbaum, Meethu Malu, Brenna McNally, Michael Muller, Leyla Norooz, Juliet Norton, Oguzhan Ozcan, Donald J. Patterson, Andreas Riener, Steven I. Ross, Karen Rust, Johannes Schnoning, M. Six Silberman, Bill Tomlinson, and Jason Yip. 2014. CHI 2039: Speculative Research Visions. *Extended Abstracts on Human Factors in Computing Systems (CHI'14)* (2014), 761–769. <https://doi.org/10.1145/2559206.2578864>
- [8] Julian Bleecker. 2009. Design Fiction: A Short Essay on Design, Science, Fact and Fiction. *Near Future Laboratory March* (2009), 4–97.
- [9] Mark Blythe. 2014. Research through design fiction. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. 703–712. <https://doi.org/10.1145/2556288.2557098>
- [10] J.C.H. Bromseth. 2002. Public places - public activities? Methodological approaches and ethical dilemmas in research on computer-mediated communication contexts. In *Reseraching ICTs in Context*, A. Morrison (Ed.). University of Oslo, 33–61.
- [11] Amy Bruckman. 2014. Research Ethics and HCI. In *Ways of Knowing in HCI*, Wendy A. Kellogg and Judith S. Olson (Eds.). Springer, 449–468.
- [12] Emanuelle Burton, Judy Goldsmith, and Nicholas Mattei. 2018. How to Teach Computer Ethics through Science Fiction. *Commun. ACM* 61, 8 (2018), 54–64. <https://doi.org/10.1145/3154485>
- [13] Stevie Chancellor, Michael L. Birnbaum, Eric D. Caine, Vincent M. B. Silenzio, and Munmun De Choudhury. 2019. A Taxonomy of Ethical Tensions in Inferring Mental Health States from Social Media. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAT*)*. Atlanta, GA.
- [14] Soumyadeb Chowdhury, Md. Sadek Ferdous, and Joemon M. Jose. 2016. Understanding Lifelog Sharing Preferences of Lifeloggers. In *Proceedings of the ACM OzCHI Conference*. <https://doi.org/10.1145/3010915.3011852>
- [15] Paul Coulton, Joseph Lindley, Miriam Sturdee, and Michael Stead. 2017. Design Fiction as World Building. *Proceedings of the Research Through Design Conference* (2017), 1–16. <https://doi.org/10.6084/m9.figshare.4746964.Image>
- [16] Kate Crawford and Megan Finn. 2015. The limits of crisis data: analytical and ethical challenges of using social and mobile data to understand disasters. *GeoJournal* 80, 4 (2015), 491–502.
- [17] Masashi Crete-Nishihata, Ronald M. Baecker, Michael Massimi, Deborah Ptak, Rachele Campigotto, Liam D. Kaufman, Adam M. Brickman, Gary R. Turner, Joshua R. Steiner, and Sandra E. Black. 2012. Reconstructing the past: Personal memory technologies are not just personal and not just for memory. *Human-Computer Interaction* 27, 1-2 (2012), 92–123. <https://doi.org/10.1080/07370024.2012.656062>
- [18] Mary Czerwinski, Douglas W. Gage, Jim Gemmell, Catherine C. Marshall, Manuel A. Perez-Quinones, Meredith M. Skeels, and Tiziana Catarci. 2006. Digital Memories in an Era of Ubiquitous Computing and Abundant Storage. *Commun. ACM* 49, 1 (2006), 45–50.
- [19] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In situ with bystanders of augmented reality glasses. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*. 2377–2386. <https://doi.org/10.1145/2556288.2557352>
- [20] Elizabeth Dubois, Anatoliy Gruz, and Jenna Jacobson. 2018. Journalists' Use of Social Media to Infer Public Opinion: The Citizens' Perspective. *Social Science Computer Review* (2018), 1–18. <https://doi.org/10.1177/0894439318791527>
- [21] James M. Dubois, Michelle Strait, and Heidi Walsh. 2017. Is It Time to Share Qualitative Research Data? *Qualitative Psychology* 5, 3 (2017), 380–393.

- [22] Brianna Dym and Casey Fiesler. 2018. Vulnerable and Online: Fandom's Case for Stronger Privacy Norms and Tools. In *Companion of the ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*. Jersey City, NJ, 329–332. <https://doi.org/10.1145/3272973.3274089>
- [23] Dave Eggers. 2013. *The Circle*. Random House, New York, NY.
- [24] Motahhare Eslami, Sneha R Krishna Kumaran, Christian Sandvig, and Karrie Karahalios. 2018. Communicating Algorithmic Process in Online Behavioral Advertising. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems* (2018), 432:1–13. <https://doi.org/10.1145/3173574.3174006>
- [25] Mylynn Felt. 2016. Social media and the social sciences: How researchers employ Big Data analytics. *Big Data & Society* 3, 1 (2016).
- [26] Md Sadek Ferdous, Soumyadeb Chowdhury, and Joemon M. Jose. 2017. Analysing privacy in visual lifelogging. *Pervasive and Mobile Computing* 40 (2017), 430–449. <https://doi.org/10.1016/j.pmcj.2017.03.003>
- [27] Casey Fiesler, Michaelanne Dye, Jessica Feuston, Chaya Hiruncharoenvate, C. J. Hutto, Shannon Morrison, Amy S. Bruckman, Munmun De Choudhury, and Eric Gilbert. 2017. What (or who) is public? Privacy settings and social media content sharing. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW)*. 567–580.
- [28] Casey Fiesler and Blake Hallinan. 2018. "We Are the Product": Public Reactions to Data Sharing and Privacy Controversies in the Media. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
- [29] Casey Fiesler and Nicholas Proferes. 2018. "Participant" Perceptions of Twitter Research Ethics. *Social Media + Society* 4, 1 (2018).
- [30] Jeana H. Frost and Michael P. Massagli. 2008. Social uses of personal health information within PatientsLikeMe, an online patient community: What can happen when patients have access to one another's data. *Journal of Medical Internet Research* 10, 3 (2008). <https://doi.org/10.2196/jmir.1053>
- [31] Sidney Fussell. 2019. You No Longer Own Your Face. *The Atlantic* (2019). <https://www.theatlantic.com/technology/archive/2019/06/universities-record-students-campuses-research/592537/>
- [32] Cathal Gurrin, Alan F. Smeaton, and Aiden R. Doherty. 2014. LifeLogging: Personal Big Data. *Foundations and Trends® in Information Retrieval* 8, 1 (2014), 1–125. <https://doi.org/10.1561/15000000033>
- [33] Blake Hallinan, Jed R Brubaker, and Casey Fiesler. 2019. Unexpected expectations: Public reaction to the Facebook emotional contagion study. *New Media & Society* (2019), 146144481987694. <https://doi.org/10.1177/1461444819876944>
- [34] Ruth Halperin and Yuval Dror. 2015. Privacy and the Digital Generation Gap: Myth and Reality. In *Proceedings of the Mediterranean Conference on Information Systems (MCIS)*. Samos, Greece, 1–14.
- [35] Brent Hecht. 2018. It's time to do something: Mitigating the negative impacts of computing through a change to the peer review process. *ACM Future of Computing Academy Blog* (2018). <https://acm-fca.org/2018/03/29/negativeimpacts/>
- [36] Anna Lauren Hoffmann and Anne Jones. 2016. Recasting Justice for Internet and Online Industry Research Ethics. In *Internet Research Ethics for the Social Age: New Cases and Challenges*, Michael Zimmer and Katharina E Kinder-kurlanda (Eds.). Peter Lang, Bern, Switzerland, 3–18.
- [37] Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014. Privacy behaviors of lifeloggers using wearable cameras. *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '14 Adjunct* (2014), 571–582. <https://doi.org/10.1145/2632048.2632079>
- [38] Jane C. Hu. 2019. Should researchers be allowed to use YouTube videos and tweets? <https://slate.com/technology/2019/06/youtube-twitter-irb-human-subjects-research-social-media-mining.html>
- [39] Xiaolei Huang, Linzi Xing, Jed R. Brubaker, and Michael J. Paul. 2017. Exploring timelines of confirmed suicide incidents through social media. *Proceedings of the IEEE International Conference on Healthcare Informatics (ICHI)* (2017), 470–477.
- [40] Luke Hutton and Tristan Henderson. 2015. "I Didn't Sign Up for This!": Informed Consent in Social Media Research. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*. 178–187. <https://research-repository.st-andrews.ac.uk/handle/10023/6691>
- [41] Krissi M. Jimroglou. 2002. A Camera With a View: JenniCAM, Visual Representation, and Cyborg Subjectivity. *Information, Communication & Society* 2, 4 (2002), 439–453. <https://doi.org/10.1080/136911899359493>
- [42] Paul Kelly, Simon J. Marshall, Hannah Badland, Jacqueline Kerr, Melody Oliver, Aiden R. Doherty, and Charlie Foster. 2013. An ethical framework for automated, wearable cameras in health behavior research. *American Journal of Preventive Medicine* 44, 3 (2013), 314–319. <https://doi.org/10.1016/j.amepre.2012.11.006>
- [43] Julie A. Kientz, Rosa I. Arriaga, and Gregory D. Abowd. 2009. Baby Steps: Evaluation of a System to Support Record-Keeping for Parents of Young Children. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. Boston, MA, 1713–1722. <https://doi.org/10.1145/1518701.1518965>
- [44] Gary King. 2011. Ensuring the data-rich future of the social sciences. *Science* 331, 6018 (2011), 719–721. <https://doi.org/10.1126/science.1197872>

- [45] Ben Kirman, John Conor Linehan, Shaun Fitz Lawson, and Dan O'Hara. 2013. CHI and the Future Robot Enslavement of Humankind; A Retrospective. In *Extended Abstracts on Human Factors in Computing Systems (CHI)*. Paris, France.
- [46] Marion Koelle, Matthias Kranz, and M Andreas. 2015. Don't look at me that way! Understanding User Attitudes Towards Data Glasses Usage. *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI)* (2015), 362–372. <https://doi.org/10.1145/2785830.2785842>
- [47] Marina Kogan, Leysia Palen, and Kenneth M. Anderson. 2016. Think local, Retweet global: Retweeting by the geographically-vulnerable during Hurricane Sandy. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW)*. 981–993.
- [48] Sze Yin Kwok, Anya Skatova, Victoria Shipp, and Andy Crabtree. 2015. The Ethical Challenges of Experience Sampling Using Wearable Cameras. In *Extended Abstracts on Mobile HCI*. Copenhagen, Denmark, 1054–1057. <https://doi.org/10.1145/2786567.2794325>
- [49] Ohbyung Kwon and Yixing Wen. 2010. An empirical study of the factors affecting social network service use. *Computers in Human Behavior* 26, 2 (2010), 254–263. <https://doi.org/10.1016/j.chb.2009.04.011>
- [50] Alex Lamb, Michael J. Paul, and Mark Dredze. 2013. Separating fact from fear: Tracking flu infections on Twitter. *HLT-NAACL* (2013), 789–795.
- [51] Matthew L Lee and Anind K Dey. 2007. Providing good memory cues for people with episodic memory impairment. *Proceedings of the International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS)* (2007), 131. <https://doi.org/10.1145/1296843.1296867>
- [52] Sam Levin. 2017. New AI can guess whether you're gay or straight from a photograph. *The Guardian* (2017). <https://www.theguardian.com/technology/2017/sep/07/new-artificial-intelligence-can-tell-whether-youre-gay-or-straight-from-a-photograph>
- [53] Joseph Lindley. 2015. Operationalising Design Fiction for Ethical Computing. *ACM SIGCAS Computers and Society* 45, 3 (2015), 79–83.
- [54] Joseph Lindley, Paul Coulton, and Miriam Sturdee. 2017. Implications for Adoption. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. Denver, CO, 265–277. <https://doi.org/10.1145/3025453.3025742>
- [55] Conor Linehan and Ben Kirman. 2014. Alternate Endings : Using Fiction. *Extended Abstracts of the ACM CHI Conference on Human Factors in Computing Systems* (2014), 45–48.
- [56] Jacob Metcalf and Kate Crawford. 2016. Where are human subjects in big data research? The emerging ethics divide. *Big Data and Society* 3, 1 (2016). <https://doi.org/10.1177/2053951716650211>
- [57] Jude Mikal, Samantha Hurst, and Mike Conway. 2016. Ethical issues in using Twitter for population-level depression monitoring: a qualitative study. *BMC Medical Ethics* 17, 1 (2016), 22. <https://doi.org/10.1186/s12910-016-0105-5>
- [58] Tze Ming Mok, Flora Cornish, and Jen Tarr. 2015. Too Much Information: Visual Research Ethics in the Age of Wearable Cameras. *Integrative Psychological and Behavioral Science* 49, 2 (2015), 309–322. <https://doi.org/10.1007/s12124-014-9289-8>
- [59] Kieron O'Hara, Mischa M. Tuffield, and Nigel Shadbolt. 2008. Lifelogging: Privacy and empowerment with memories for life. *Identity in the Information Society* 1, 1 (2008), 155–172. <https://doi.org/10.1007/s12394-009-0008-4>
- [60] Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. 2010. Earthquake shakes Twitter users: Real-time event detection by social sensors. *Proceedings of the International Conference on World Wide Web (WWW)* (2010), 851–860.
- [61] Sara Shaw and Geraldine Barrett. 2006. Research governance: Regulating risk and reducing harm? *Journal of the Royal Society of Medicine* 99, 1 (2006), 14–19. <https://doi.org/10.1258/jrsm.99.1.14>
- [62] A. Signorini, A.M. Segre, and P.M. Polgreen. 2011. The use of Twitter to track levels of disease activity and public concern in the US during the influenza A H1N1 pandemic. *PLoS ONE* 6, 5 (2011).
- [63] Natasha Singer. 2015. Love in the Time of Twitter. <https://bits.blogs.nytimes.com/2015/02/13/love-in-the-times-of-twitter/>
- [64] Robert Soden, Michael Skirpan, Casey Fiesler, Zahra Ashktorab, Eric Baumer, Mark Blythe, and Jasmine Jones. 2019. CHI4EVIL: Creative Speculation On the Negative Impacts of HCI Research. In *Extended Abstracts on Human Factors in Computing Systems (CHI EA)*.
- [65] Luke Stark. 2018. Algorithmic psychometrics and the scalable subject. *Social Studies of Science* 48, 2 (2018), 204–231. <https://doi.org/10.1177/0306312718772094>
- [66] United States. 1978. *The Belmont report: Ethical principles and guidelines for the protection of human subjects of research*. Technical Report. The Commission, Bethesda, MD.
- [67] Fred Stutzman, Ralph Gross, and Alessandro Acquisti. 2013. Silent Listeners: The Evolution of Privacy and Disclosure on Facebook. *Journal of Privacy and Confidentiality* 4, 2 (2013), 7–41. <http://repository.cmu.edu/jpc/vol4/iss2/2/>
- [68] Robert Templeman, Mohammed Korayem, David Crandall, and Apu Kapadia. 2014. PlaceAvoider: Steering First-Person Cameras away from Sensitive Spaces. *Proceedings of the Network and Distributed System Security (NDSS) Symposium* February (2014), 23–26. <https://doi.org/10.14722/ndss.2014.23014>

- [69] Edison Thomaz, Aman Parnami, Jonathan Bidwell, Irfan Essa, and Gregory D. Abowd. 2013. Technological approaches for addressing privacy concerns when recognizing eating behaviors with wearable cameras. *Proceedings of the 2013 ACM International Conference on Pervasive and Ubiquitous Computing (UbiComp)* (2013), 739. <https://doi.org/10.1145/2493432.2493509>
- [70] Elizabeth Tucker. 2014. The Endangered Child: Choking and Fainting Games in the Online Underground of YouTube. *Children's Folklore Review* 36 (2014), 19–34.
- [71] 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. <https://doi.org/10.1177/1556264617725200>
- [72] Jessica Vitak, Katie Shilton, and Z. Ashktorab. 2016. Beyond the Belmont Principles: Ethical Challenges, Practices, and Beliefs in the Online Data Research Community. In *Proceedings of the ACM CSCW Conference on Computer Supported Cooperative Work & Social Computing*. San Francisco, CA, 941–953.
- [73] Helena Webb, Marina Jirotko, Bernd Carsten Stahl, William Housley, Adam Edwards, Matthew Williams, Rob Procter, Omer Rana, and Pete Burnap. 2017. The Ethical Challenges of Publishing Twitter Data for Research Dissemination. *Proceedings of the ACM on Web Science Conference (WebSci)* (2017), 339–348. <https://doi.org/10.1145/3091478.3091489>
- [74] Daniel Weiz, Gagat Anad, and Claus-Peter H. Ernst. 2016. The Influence of Subjective Norm on the Usage of Smartglasses. In *The Drivers of Wearable Device Usage*, Claus-Peter H. Ernst (Ed.). Springer, 1–11.
- [75] Katrin Weller and Katharina E Kinder-kurlanda. 2015. Uncovering the Challenges in Collection , Sharing and Documentation : The Hidden Data of Social Media Research? *Proceedings of the Ninth International AAAI Conference on Web and Social Media* 2014 (2015), 28–37.
- [76] Lindy West. 2015. What happened when I confronted my cruelest troll. *The Guardian* (2015). <https://www.theguardian.com/society/2015/feb/02/what-happened-confronted-cruellest-troll-lindy-west>
- [77] 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), 111:1–27. <https://doi.org/10.1145/3134746>
- [78] Richmond Y. Wong, Ellen Van Wyk, and James Pierce. 2017. Real-fictional entanglements: Using science fiction and design fiction to interrogate sensing technologies. *Proceedings of the ACM DIS Conference on Designing Interactive Systems* (2017), 567–579. <https://doi.org/10.1145/3064663.3064682>
- [79] Sean D. Young, Caitlin Rivers, and Bryan Lewis. 2014. Methods of using real-time social media technologies for detection and remote monitoring of HIV outcomes. *Preventive Medicine* 63 (2014), 112–115.
- [80] Michael Zimmer. 2010. "But the data is already public": On the ethics of research in Facebook. *Ethics and Information Technology* 12, 4 (2010), 313–325. <https://doi.org/10.1007/s10676-010-9227-5>
- [81] Michael Zimmer. 2016. OKCupid Study Reveals the Perils of Big Data Science. <https://www.wired.com/2016/05/okcupid-study-reveals-perils-big-data-science/>
- [82] Michael Zimmer and Nicholas John Proferes. 2014. A Topology of Twitter Research: Disciplines, Methods, and Ethics. *Aslib Journal of Information Management* 66, 3 (2014), 250–261.

Received June 2019; revised August 2019; accepted October 2019