

Business (mis)Use Cases of Generative AI

Stephanie Houde, Vera Liao, Jacquelyn Martino, Michael Muller

David Piorkowski, John Richards, Justin Weisz, Yunfeng Zhang

Stephanie.Houde@ibm.com, Vera.Liao@ibm.com, jmartino@us.ibm.com, michael_muller@us.ibm.com

David.Piorkowski@ibm.com, ajtr@us.ibm.com, jweisz@us.ibm.com, zhangyun@us.ibm.com

IBM Research

ABSTRACT

Generative AI is a class of machine learning technology that learns to generate new data from training data. While deep fakes and media-and art-related generative AI breakthroughs have recently caught people’s attention and imagination, the overall area is in its infancy for business use. Further, little is known about generative AI’s potential for malicious misuse at large scale. Using co-creation design fictions with AI engineers, we explore the plausibility and severity of business misuse cases.

CCS CONCEPTS

• **Human-centered computing**; • **Social and professional topics** → *Computing / technology policy*; • **Applied computing** → *Law, social and behavioral sciences*;

KEYWORDS

Generative AI, design fiction, human-in-the-loop

ACM Reference Format:

Stephanie Houde, Vera Liao, Jacquelyn Martino, Michael Muller and David Piorkowski, John Richards, Justin Weisz, Yunfeng Zhang. 2020. Business (mis)Use Cases of Generative AI. In *Proceedings of March 17, 2020 (IUI '20 Workshops)*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

As we think about a future where humans and AI partner in creative activities, we consider how generative models [40, 41, 51, 53, 60] could impact current businesses and possibly create new ones. While deep fakes and media-and art-related generative AI breakthroughs have recently caught people’s attention and imagination [2, 23, 27, 28, 45], the area overall is in its infancy for business use. In this paper, we take an inverse approach to business cases and propose business “misuse” cases from the “bad actor” point of view. Using a practice called design fictions [8, 18, 19, 48–50], we engaged with software engineers who are expert in AI technologies. We provided three half-page fictions about possible harmful applications of generative technologies as probes for use in our co-creation exercises. As we guided the engineers’ exploration of the three fabricated misuse business cases, we learned their points

of view on factors such as scenario plausibility, seriousness, and prevention.

The contributions of this paper include

- participant-created future scenarios based on generative AI capabilities in text, audio, and video
- reactions to possible generative AI design fiction scenarios
- discussion responses to specific probes on scenario plausibility, seriousness, ways things might be worse, ways things might be better, and prevention

The remainder of this paper is organized as follows: Section 2 discusses related work in AI and Human Centered Data Science (HCDS), and then introduces design fictions as a method for “prototyping” and discussing possible future outcomes of AI. Section 3 describes how we used design fictions with experts in AI software engineering to explore potential business misuse cases. Section 4 presents our results. Section 5 offers concluding thoughts.

2 RELATED WORK

Strong claims are made about the promises and current successes of AI and data science [20, 25, 32, 52]. While some of these claims are projected for the future [16], Agarwal and Dhar editorialized in 2014 that “This is powerful... we are in principle already there” [1]. Meanwhile to the dystopian extreme, scholars warn about the “mythology” of working with big data, that the quantitative nature of data gives a false illusion that all data-driven outcomes are objective, ethical, or true [12]. Increasingly, scholars in the emerging field of Human Centered Data Science (HCDS) [3, 11, 33, 55] have begun to investigate data science practices, showing the necessary, responsible, and increasingly accountable human activities that take place between data and models [11, 21, 22, 26, 34, 37–39, 42].

In this paper we extend this emerging work by examining possible applications of recent developments in generative AI methods. Rather than waiting to “see what happens,” we apply the low-cost method of design fictions to prototype [6, 18, 19, 30, 31, 36, 46–50, 54, 56] possible future applications. Rather than depending on our own views, we engage with the thoughtful and creative contributions of knowledgeable colleagues [4, 7, 8, 14, 29, 35, 57].

There has been increasing interest in HCI of projecting the future of technologies through design fictions (DFs) – fictional scenarios in the form of narratives, concepts, prototypes, enactments, and games. [7–9, 14]. DFs have been applied to the design of new technology [15], to explore how future users may adopt a technology [14], and as critical tools to anticipate the social and political consequences of technologies [17, 49].

Design fictions as a methodology spring from several sources. While many cite Sterling’s introductory definition of “deliberate use of diegetic prototypes to suspend disbelief about change” [48], relevant forms have also been explored in the fields of Participatory

Design [5, 10], and Future Studies [36]. A predecessor to DFs in HCI research is design scenarios [59] - simple vignettes to illustrate the use of technologies to be developed. DFs sometimes focus on the design of concrete “near future” technologies [6], constructing a discursive fictional space for speculating about the effects of such technologies in the future. In particular, DFs has been embraced for revealing values associated with new technologies [35, 49].

There are many forms and usages of design fiction. We differentiate between research works that create DFs as the end products (e.g., to frame new design concepts that do not yet exist [13, 24]), or as critical tools [7, 8]), and those using DFs as probes—“critical narratives to elicit open-ended responses from potential future users [stakeholders] of proposed technologies” [43].

Our use of DFs falls in the latter category. Some also refer to it as participatory design fiction, in contrast to previous approaches where DFs are created only by researchers [35]. For example, Shulte et al. described a 5-step method to create design fictions on the topic of smart houses, and illustrated how the method can be used for research and design purposes. Recently, several works explored using the Story Completion Method (SCM) in design fiction. SCM was first introduced in psychotherapy and qualitative research in psychology [44]. Wood et al. used SCM for speculative stories to explore the future vision of Virtual Reality pornography [58]. Cheon and Su introduced Futuristic Autobiographies [14] as a way to elicit perspectives on the future of technologies and conducted a case study on the future of robotics. The method starts by posing stories involving the participant as a character in a future state, as grounded in background research work, and invites the participant to complete the autobiography.

3 METHODOLOGY

Our exploratory study into potential business misuse cases began with the creation of design fictions focusing on the use of generative AI for malicious purposes. We wrote three scenarios, each representing a different generative media type: text, audio, or video. Then, using these scenarios as a springboard, we led co-creation sessions with AI experts to gain their views on possible worst-case business scenarios. We conducted these sessions in counterbalanced order across participants, to control for novelty, fatigue, and other possible order effects.

3.1 Interviews

We conducted co-creation sessions with six participants (2 female, 4 male) from our research organization. All participants are software engineers practiced in AI techniques. After a level-setting description of generative techniques as a class of machine learning and their potential use for malicious intent, we presented our potential future scenarios. Each scenario was presented in two parts. The first part is a design fiction from the near future (2020). After reading the first part, we asked each participant to speculate on a worst case misuse of generative AI in ten more years (2030) and invent their own future scenario. After hearing their thoughts, we shared the second part which was our version of the future in 2030. With both our version and their version in mind, we asked them to respond to the following questions:

(1) Is the scenario plausible?

(2) Is this a serious problem?

(3) Can you make the scenario worse?

(4) Can you make the scenario better?

(5) Can anything be done to prevent this?

(6) Is there a way out of this?

(7) Who is the right person, organization, or entity to improve it?

Participants spent about 15 minutes reading and discussing each scenario for a total time of about 45 minutes.

3.2 Scenarios

3.2.1 Scenario A. Instant Author Story.

Part 1. It’s 2020 (and you’ve just seen this popup ad): Hey! I bet you’ve tried to write the Great American Novel several times but you can’t seem to get past the first chapter. We get it. Like you, we find original writing to be really difficult and basically unrewarding. Luckily, we’ve created MakeMyStory.com that takes someone else’s content, does a deep analysis of it, and emits it in a form that is as compelling as the original work but different enough that no one will say it’s not yours. How cool is that? As an added bonus, your new work can be emitted in any language, increasing the number of countries where it can be sold.

Part 2. It’s 2030 now and Amazon is completely flooded with works by previously unknown writers. A few prominent authors have taken MakeMyStory to court to try to have the web site shut down. Others have sued the nominal authors of these works but have not been successful. This is due to a few things. First, MakeMyStory has declared their software to be proprietary so it can’t be examined and the courts have gone along with that so far (go figure). Second, the work that is claimed to be derived passes every known test for originality. (Interesting side note: perhaps an investigation into how the original authors determined that their work had been stolen might lead to new techniques for automatically finding such “adaptations”.) Third, a recently-added REMIX capability allows multiple original works to be “blended” making derivation tracking almost impossible. So ... check back in in another couple of years, assuming anyone is still bothering to read by then.

3.2.2 Scenario B. Fake Smoking Evidence Insurance Claim Story.

Part 1. It’s 2020 and Jane just accepted a new job. As part of her onboarding process, she is answering some health questions for insurance purposes, including, of course, whether she smokes. That one is simple...never. Later that year, Jane, develops shortness of breath, confirmed to be the early stages of emphysema. Confident that her health insurance will cover the costs she begins some very expensive treatment. When the first batch of bills is rejected by the insurance company she calls to find out why and is simply shocked. Here is the conversation:

Representative: I’m sorry Jane, but based on your medical history, you will not be covered for this procedure. Your history of smoking disqualifies you from coverage.

Jane: What are you talking about, I’ve never smoked in my life!

Representative: Well, I’m looking at some videos online that beg to differ. Did you attend a wedding last month?

Jane: Yes, it was my best friend’s wedding. But I don’t understand... what are you talking about?

Representative: You were wearing a blue dress, correct?

Jane: Yes.

Representative: Well ma'am, I'm watching a video from the wedding where you are very clearly smoking a cigarette. As such your request for coverage has been denied.

Jane knows that video had to have been faked, but she has no idea how to prove it. Let's just hope her savings hold out.

Part 2. In what seems to be a pattern of corruption, insurance companies are denying medical claims due to video "evidence" of risky behavior ranging from base jumping to travel to active war zones. Car insurance claims are being rejected due to "evidence" of preexisting damage. Even life insurance claims are being denied based on "evidence" that the deceased is still alive. Legislation was passed to make these practices illegal, but without the ability to detect when a video is faked, and without the ability to truly prove the legitimacy of a piece of digital content, insurance companies are able to continue these practices to fatten their bottom line.

3.2.3 Scenario C. Audio Evidence Story.

Part 1. It's 2020, and the following story just appeared on page 14 of the New York Times: Expert testimony for the defense was presented today in the case of Walter Milgrim, a prominent New York investor accused of arranging the contract murder of his wife. In earlier testimony, the prosecution's expert witness had played an audio recording claimed to have been captured by a passenger on the Chappaqua train platform on the night of the murder. In it, a person sounding like Milgrim and a person sounding like the trigger man could be heard discussing the means of payment. While the voices were somewhat muffled, the recording seemed authentic. This was cleverly refuted by the defense expert, who commissioned a separate recording at the train station. In that recording, a low hum was heard due to wind passing through the overhead wires. This hum was not present in the prosecution's recording, but it was to be expected due to windy conditions that night. Reasonable doubt? We'll see.

Part 2. It's 2030 and a number of criminal convictions have been thrown into turmoil by the discovery of an audio tool that can synthesize ultra-realistic ambient soundscapes based on thousands of parameters. For example, all manner of weather conditions can be simulated including how those conditions interact with physical features of a location. Seemingly-authentic crowd noises (from the faint shuffle of feet on snow to the sounds of people almost having to shout over traffic or being nearly drowned out by the arrival of a train or airplane) add to the realism. A whistleblower has alleged that a number of high profile cases in the Southern District of New York relied on recordings generated by this tool. It is unknown how many cases have been corrupted and there seems to be no way to determine which recordings were real and which were, in fact, generated.

4 RESULTS

Two members of the research team independently reviewed the transcripts and audio recordings of each of the three scenarios in each of the six sessions, looking for key themes and comments. These were then assembled into a combined data set. As no clear differences between the three scenarios were found, we present observations summarized by themes within session phases, starting

with participant comments during their invented 2030 futures, followed by their reactions to our envisioned 2030 futures, and ending with their responses to the focused questions.

4.1 Participant Invented Futures

4.1.1 Ubiquity. Participants anticipated a significant increase in fakes due to reduced costs. Many of these fakes would likely be created by the general public. It was noted that large scale automated faking might be particularly dangerous since it would spread adverse impact beyond what a few bad actors would likely do. Interestingly, the very ubiquity of fakes might also reduce the impact of any particular fake since less weight would be given to any single piece of content. Widespread surveillance, with verifiable provenance, might become more important as a means of countering faked evidence, and privacy laws might require revisions as policy-makers learn more about how vulnerable groups could be affected by the proliferation of fakes, and in what ways. Not surprisingly, there will be increased need for reliable detection of fakes.

4.1.2 Quality. When informants were discussing *their own participatory fictions*, they suggested that machine generated text may actually be better than human authored text, at least for some purposes. It was also noted that faked content may actually be acceptable if it is used for entertainment rather than as evidence.

4.1.3 Ownership. There was a belief that all work was, in some sense, derived from what came before. This view led to the thought that automatically derived content was perhaps acceptable.

4.1.4 Markets. Niche markets for hand-crafted content might emerge with "written by humans" sections being tucked in the back of book stores and online sales sites. A greatly expanded need for digital forensics in cases of disputed ownership and other matters of evidence might cause this to become an emerging business opportunity.

4.1.5 Technology. It was felt that models may be able to detect features of fakes that would not be apparent to humans.

4.2 Reactions To Our Envisioned Futures

4.2.1 Ubiquity. People might start creating verified videos of themselves to protect against fakes. This is already happening with always-on dash cams in both private and police cars. Because of the ease of fake creation, multiple independent sources of evidence might become even more required.

4.2.2 Quality. When informants were discussing *our probe fictions*, they thought that there was a real possibility of a descent into mediocrity as fakes fed on themselves in a "never ending resounding echo chamber of cyber-generated nonsense". Several thought we risked a loss of "art", both in the creation of and the appreciation of content, with creatives having significantly reduced value. It is possible this would not be sustained, though, as much generated content might be "boring as hell".

4.2.3 Ownership. One participant wondered if AI would have to be acknowledged as a co-creator of content.

4.2.4 Markets. Creating content for a single consumer would be feasible, both for entertainment and education. This would upend the current mass distribution market model.

4.2.5 Technology. It was thought that there would be a significant ramp up of research in detecting machine-generated content. Progress here would likely be aided by the fact that detection capability would increase as the corpus of fakes for training grew. On the other hand, widespread watermarking and the use of block chains might obviate the need for much of this detection capability as we could rely on immutable metadata to prove originality.

4.3 Answers To Focused Questions

4.3.1 Plausibility. There were mixed views on the plausibility of 2030 scenarios. Some felt that similar futures were inevitable due to reduced costs of fakes and often malevolent human nature. Others considered these futures unlikely due to legal, regulatory, and market forces, the emergence of technical detection and provenance mechanisms, and the belief that societies have generally been good at authenticating things they care about (e.g., money, passports).

4.3.2 Seriousness. There were also mixed views on the seriousness of these possible futures; if a future was considered implausible for the reasons noted above it was also viewed as not serious. On the other hand, it was mentioned that a single blemish in a person's health record could profoundly impact the ability to obtain insurance coverage. Thus, quite serious outcomes might arise from even isolated incidents of faking.

4.3.3 How It Could Be Worse. One participant felt that generative models interactively posing as humans might be particularly harmful. We already see low-tech variants of this with, for example, people posing as somewhat-distant relatives needing money to get them out of difficulty during foreign travel. The synthesis of realistic human speech was considered harder than the synthesis of ambient audio, but the potential impact could be much worse. On a different note, the coherent faking of multiple sources might further complicate use as evidence. One possible reaction to all this would be governments banning the use of generative AI. If *only* governments were allowed to use it, the results could be quite dire. Finally, if these sorts of futures come to pass, people might simply stop caring about truth altogether.

4.3.4 How It Could Be Better. Several potentially beneficial uses of generative AI were offered. In one, the AI could generate realistic time-aged portraits of lost loved ones to convey the sense that they were still present in peoples' lives. In another, high-quality generated speech could be used by those with severe communication disabilities. Related to this, a generative AI might align the sounds of diverse human accents in an audio stream for improved mutual intelligibility when time was critical (e.g., in an emergency call). Finally, the ability to cheaply produce content could provide value to currently under-served markets.

4.3.5 How It Could Be Prevented. Answers to the "Prevention" and "Way Out" questions were quite similar and are combined here. Most felt that some combination of technology, regulation, and legislation could prevent the worst of these future scenarios. Data traceability and auditability, ubiquitous detection of fake content

and watermarking of captured content in personal devices, and new legislation that finally caught up with emerging threats were all mentioned. Education and training were noted by many participants as important; making people aware of the threats posed by fake content and equipping them with skills to detect it might diminish the viral quality of fakes.

4.3.6 Who Might Improve. Here too, some combination of technology providers, regulators (especially international organizations), and legislative bodies were viewed as likely contributors. Specialized oversight mechanisms might also arise to meet the needs of academic research and publishing, medical research, art competitions, and so on. One participant thought that AI itself might assist in doing legal research and crafting good law to counter threats. Finally, consumers and critics might rate and review content for quality and originality in a manner not too dissimilar from today.

5 CONCLUDING THOUGHTS

Our participants' views were both thoughtful and highly varied. Even though their comments stayed at a quite general level, they sometimes came to contrasting conclusions when discussing their own future scenarios as opposed to the future scenarios that we created (particularly regarding quality issues). These types of contrasts in perspectives have been important in other projects with participatory forms of design fictions [4, 7, 8, 14, 29, 35, 57]. Taken together, these contrasting results suggest that *including informants in collaborative interpretive spaces* can be a powerful method for increasing the knowledge of participants and researchers alike.

Some informants thought scenarios of the sort we sketched for 2030 were inevitable while others thought they were unlikely. Drawing on their own expertise, their thoughts on how technology might counter possible threats were more nuanced than their thoughts about regulation and legislation. Perhaps the most interesting dynamics arose in connection with what many saw as an "arms race" between generation and detection capabilities. There was optimism that the coming ubiquity of fakes will lead, in a manner not unlike a natural ecosystem, to an abundance of training data for high-quality detectors. That, along with progress in digital provenance and updated regulations and legislation, may prevent the worst of these future scenarios from coming to pass. It will be important to expand these HCDS approaches [3, 11, 33, 55] to include more diverse informants in the future, for a more societally-grounded vision and critique of the implications of generative AI.

REFERENCES

- [1] Ritu Agarwal and Vasant Dhar. 2014. Big data, data science, and analytics: The opportunity and challenge for IS research.
- [2] Sakshi Agarwal and Lav R Varshney. 2019. Limits of Deepfake Detection: A Robust Estimation Viewpoint. *arXiv preprint arXiv:1905.03493* (2019).
- [3] Cecilia Aragon, Clayton Hutto, Andy Echenique, Brittany Fiore-Gartland, Yun Huang, Jinyoung Kim, Gina Neff, Wanli Xing, and Joseph Bayer. 2016. Developing a research agenda for human-centered data science. In *Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion*. ACM, 529–535.
- [4] Eric PS Baumer, June Ahn, Mei Bie, Elizabeth M Bonsignore, Ahmet Börütece, Oğuz Turan Buruk, Tamara Clegg, Allison Druin, Florian Ehtler, Dan Gruen, et al. 2014. CHI 2039: speculative research visions. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*. ACM, 761–770.
- [5] Ian Beeson and Clodagh Miskelly. 2000. Dialogue and dissent in stories of community. In *Proc. PDC 2000*. Citeseer.

- [6] J Bleecker. 2009. Design Fiction: A Short Essay. Available at <http://nearfuturelaboratory.com/2009/03/17/design-fiction-a-short-essay-on-design-science-fact-andfiction/> [date accessed: 6th Jan 2017] (2009).
- [7] Mark Blythe and Elizabeth Buie. 2014. Chatbots of the gods: imaginary abstracts for techno-spirituality research. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational*. ACM, 227–236.
- [8] Mark Blythe, Enrique Encinas, Jofish Kaye, Miriam Lueck Avery, Rob McCabe, and Kristina Andersen. 2018. Imaginary design workbooks: Constructive criticism and practical provocation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 233.
- [9] Mark Blythe, Jamie Steane, Jenny Roe, and Caroline Oliver. 2015. Solutionism, the game: design fictions for positive aging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 3849–3858.
- [10] Susanne Bødker and Kaj Grønbaek. 1991. Design in action: From prototyping by demonstration to cooperative prototyping. In *Design at work: Cooperative design of computer systems*. Lawrence Erlbaum Associates, 197–218.
- [11] Nadia Boukhelifa, Anastasia Bezerianos, Ioan Cristian Trelea, Nathalie Méjean Perrot, and Evelyne Lutton. 2019. An Exploratory Study on Visual Exploration of Model Simulations by Multiple Types of Experts. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 644.
- [12] Danah Boyd and Kate Crawford. 2012. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, communication & society* 15, 5 (2012), 662–679.
- [13] Barry Brown, Julian Bleecker, Marco D'adamo, Pedro Ferreira, Joakim Formo, Mareike Glöss, Maria Holm, Kristina Höök, Eva-Carin Banka Johnson, Emil Kaburuan, et al. 2016. The IKEA Catalogue: Design fiction in academic and industrial collaborations. In *Proceedings of the 19th International Conference on Supporting Group Work*. ACM, 335–344.
- [14] EunJeong Cheon and Norman Makoto Su. 2018. Futuristic autobiographies: Weaving participant narratives to elicit values around robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 388–397.
- [15] Nazli Cila, Iskander Smit, Elisa Giaccardi, and Ben Kröse. 2017. Products as agents: metaphors for designing the products of the IoT age. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 448–459.
- [16] Ciprian Dobre and Fatos Xhafa. 2014. Intelligent services for big data science. *Future Generation Computer Systems* 37 (2014), 267–281.
- [17] Paul Dourish and Genevieve Bell. 2014. Resistance is futile: reading science fiction alongside ubiquitous computing. *Personal and Ubiquitous Computing* 18, 4 (2014), 769–778.
- [18] Anthony Dunne and Fiona Raby. 2001. *Design noir: The secret life of electronic objects*. Springer Science & Business Media.
- [19] Anthony Dunne and Fiona Raby. 2013. *Speculative everything: design, fiction, and social dreaming*. MIT press.
- [20] Hugh Durrant-Whyte. 2015. Data, knowledge and discovery: Machine learning meets natural science. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 7–7.
- [21] Melanie Feinberg. 2017. A design perspective on data. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2952–2963.
- [22] Lisa Gitelman. 2013. *Raw data is an oxymoron*. MIT press.
- [23] David Güera and Edward J Delp. 2018. Deepfake video detection using recurrent neural networks. In *2018 15th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. IEEE, 1–6.
- [24] Felix Heibeck, Alexis Hope, and Julie Legault. 2014. Sensory fiction: a design fiction of emotional computation. In *Proceedings of the 2nd ACM International Workshop on Immersive Media Experiences*. ACM, 35–40.
- [25] Anthony JG Hey, Stewart Tansley, Kristin M Tolle, et al. 2009. *The fourth paradigm: data-intensive scientific discovery*. Vol. 1. Microsoft research Redmond, WA.
- [26] Sean Kandel, Andreas Paepcke, Joseph Hellerstein, and Jeffrey Heer. 2011. Wrangler: Interactive visual specification of data transformation scripts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3363–3372.
- [27] Pavel Korshunov and Sébastien Marcel. 2019. Vulnerability assessment and detection of deepfake videos. In *The 12th IAPR International Conference on Biometrics (ICB)*. 1–6.
- [28] Yuezun Li and Siwei Lyu. 2018. Exposing deepfake videos by detecting face warping artifacts. *arXiv preprint arXiv:1811.00656* 2 (2018).
- [29] Ann Light. 2011. Democratizing technology: making transformation using designing, performance and props. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2239–2242.
- [30] Joseph Lindley. 2015. Researching Design Fiction With Design Fiction. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*. ACM, 325–326.
- [31] Joseph Lindley and Robert Potts. 2014. A machine learning: an example of HCI prototyping with design fiction. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational*. ACM, 1081–1084.
- [32] Renée J Miller. 2017. The future of data integration. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 3–3.
- [33] Michael Muller, Melanie Feinberg, Timothy George, Steven J Jackson, Bonnie E John, Mary Beth Kery, and Samir Passi. 2019. Human-Centered Study of Data Science Work Practices. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, W15.
- [34] Michael Muller, Ingrid Lange, Dakuo Wang, David Piorkowski, Jason Tsay, Q Vera Liao, Casey Dugan, and Thomas Erickson. 2019. How Data Science Workers Work with Data: Discovery, Capture, Curation, Design, Creation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 126.
- [35] Michael Muller and Q Vera Liao. 2017. Exploring AI Ethics and Values through Participatory Design Fictions. *Human Computer Interaction Consortium* (2017).
- [36] Daniel Pargman, Elina Eriksson, Mattias Höjer, Ulrika Gunnarsson Östling, and Luciane Aguiar Borges. 2017. The (un) sustainability of imagined future information societies. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 773–785.
- [37] Samir Passi and Steven Jackson. 2017. Data vision: Learning to see through algorithmic abstraction. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 2436–2447.
- [38] Samir Passi and Steven J Jackson. 2018. Trust in Data Science: Collaboration, Translation, and Accountability in Corporate Data Science Projects. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 136.
- [39] Kathleen H Pine and Max Liboiron. 2015. The politics of measurement and action. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 3147–3156.
- [40] Alec Radford, Luke Metz, and Soumith Chintala. 2015. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv e-prints*, Article arXiv:1511.06434 (Nov 2015), arXiv:1511.06434 pages. arXiv:cs.LG/1511.06434
- [41] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. (2019).
- [42] Tye Rattenbury, Joseph M Hellerstein, Jeffrey Heer, Sean Kandel, and Connor Carvers. 2017. *Principles of data wrangling: Practical techniques for data preparation*. "O'Reilly Media, Inc."
- [43] Britta F Schulte, Paul Marshall, and Anna L Cox. 2016. Homes for life: a design fiction probe. In *Proceedings of the 9th nordic conference on human-computer interaction*. ACM, 80.
- [44] Iduna Shah-Beckley, Victoria Clarke, and Zoe Thomas. 2018. Therapists' and non-therapists' constructions of heterosex: A qualitative story completion study. *Psychology and Psychotherapy: Theory, Research and Practice* (2018).
- [45] Yevhenii SHTEFANIUK, Ivan OPIRSKYI, Ihor IVANCHENKO, and Kateryna PINDEL. [n.d.]. DEEPFAKE-NEW TECHNOLOGY FOR IMPERSONATION IN CYBERATTACKS. ([n.d.]).
- [46] Michael Skirpan and Casey Fiesler. 2018. Ad empathy: A design fiction. In *Proceedings of the 2018 ACM Conference on Supporting Groupwork*. ACM, 267–273.
- [47] Marie Louise Juul Søndergaard and Lone Koefoed Hansen. 2016. PeriodShare: A Bloody Design Fiction. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. ACM, 113.
- [48] Bruce Sterling. 2009. Design fiction. *interactions* 16, 3 (2009), 20–24.
- [49] Joshua Tanenbaum. 2014. Design fictional interactions: why HCI should care about stories. *interactions* 21, 5 (2014), 22–23.
- [50] Joshua Tanenbaum, Marcel Pufal, and Karen Tanenbaum. 2016. The limits of our imagination: design fiction as a strategy for engaging with dystopian futures. In *Proceedings of the Second Workshop on Computing within Limits*. ACM, 10.
- [51] V Tsourikov. 2017. Generative inventions. In *TRIZfest, Conference Proceedings*. 178–184.
- [52] Wil MP Van der Aalst. 2014. Data scientist: The engineer of the future. In *Enterprise interoperability VI*. Springer, 13–26.
- [53] Tijn van der Zant, Matthijs Kouw, and Lambert Schomaker. 2013. Generative artificial intelligence. In *Philosophy and Theory of Artificial Intelligence*. Springer, 107–120.
- [54] Ron Wakkary, William Odom, Sabrina Hauser, Garnet Hertz, and Henry Lin. 2015. Material speculation: actual artifacts for critical inquiry. In *Proceedings of The Fifth Decennial Aarhus Conference on Critical Alternatives*. Aarhus University Press, 97–108.
- [55] Dakuo Wang, Justin D Weisz, Michael Muller, Parikshit Ram, Werner Geyer, Casey Dugan, Yla Tausczik, Horst Samulowitz, and Alexander Gray. 2019. Human-AI Collaboration in Data Science: Exploring Data Scientists' Perceptions of Automated AI. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 211.
- [56] Richmond Y Wong and Deirdre K Mulligan. 2016. When a product is still fictional: anticipating and speculating futures through concept videos. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. ACM, 121–133.
- [57] 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.

- [58] Matthew Wood, Gavin Wood, and Madeline Balaam. 2017. They're Just Tixel Pits, Man: Disputing the 'Reality' of Virtual Reality Pornography through the Story Completion Method. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 5439–5451.
- [59] David Wright, Rachel Finn, Raphael Gellert, Serge Gutwirth, Philip Schütz, Michael Friedewald, Silvia Venier, and Emilio Mordini. 2014. Ethical dilemma scenarios and emerging technologies. *Technological Forecasting and Social Change* 87 (2014), 325–336.
- [60] Jonathan L Zittrain. 2005. The generative internet. *Harv. L. Rev.* 119 (2005), 1975.