



Designing Technology and Policy Simultaneously: Towards A Research Agenda and New Practice

Qian Yang
Cornell University
USA
qianyang@cornell.edu

Richmond Wong
Georgia Institute of
Technology
USA
rwong34@gatech.edu

Thomas Gilbert
Cornell Tech
USA
tj299@cornell.edu

Margaret Hagan
Stanford University
USA
mdhagan@stanford.edu

Steven Jackson
Cornell University
USA
sjj54@cornell.edu

Sabine Junginger
The Lucerne University of
Applied Sciences and Arts
Switzerland
sabine.junginger@hslu.ch

John Zimmerman
Carnegie Mellon University
USA
johnz@cs.cmu.edu

ABSTRACT

Accounting for technologies' unintended consequences—whether they are misinformation on social media or issues of sustainability and social justice—increasingly requires HCI to consider technology design at a societal-level scale. At this scale, public and corporate policies play a critical role in shaping technologies and user behaviors. However, the research and practices around tech and policy design have largely been held separate. How can technology design and policies better inform and coordinate with each other in generating safe new technologies? What new solutions might emerge when HCI practitioners design technology and its policies simultaneously to account for its societal impacts? This workshop addresses these questions. It will 1) identify disciplines and areas of expertise needed for a tighter, more proactive technology-and-policy-design integration, 2) launch a community of researchers, educators, and designers interested in this integration, 3) identify and publish an HCI research and education agenda towards designing technologies and technology policies simultaneously.

CCS CONCEPTS

- **Human-centered computing** → *Collaborative and social computing design and evaluation methods; HCI design and evaluation methods;*
- **Social and professional topics** → *Computing / technology policy.*

ACM Reference Format:

Qian Yang, Richmond Wong, Thomas Gilbert, Margaret Hagan, Steven Jackson, Sabine Junginger, and John Zimmerman. 2023. Designing Technology and Policy Simultaneously: Towards A Research Agenda and New Practice. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23), April 23–28, 2023, Hamburg, Germany*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3544549.3573827>

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

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

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

ACM ISBN 978-1-4503-9422-2/23/04.

<https://doi.org/10.1145/3544549.3573827>

1 MOTIVATION & WORKSHOP GOALS

Accounting for technology's unintended consequences constitutes an important and growing HCI concern. From the spread of misinformation on social media to privacy concerns looming over smart homes and smart cities, technological perils today often result from a complex interplay among technologies, user behaviors, and the economic, legal, and social contexts that undergird all. Preventing such perils requires not only discrete, well-intentioned technology designs, but also effective public and corporate policies that govern technologies and user behaviors at large and over time [14, 15, 19]. Consider web accessibility as an example. Accessible web designs in HCI informed policies both at national (e.g., Americans with Disabilities Act Standards for Accessible Design [1]) and corporate levels (e.g., Twitter image alt text requirements [17].) In turn, these policies helped disseminate accessible web designs and made them more durable by enforcing legal responsibilities [15, 16]. Through different mechanisms, tech designs and policies can complement each other in promoting new technologies while preventing their unintended consequences.

Prior HCI research has argued that technology designs and policies are not only complementary; they are interdependent and should be considered simultaneously [10]. Public and corporate policies pre-configure technology design, closing some design possibilities while opening up others. Including policy design in technology design processes can explicate and even challenge the economic, legal, and organizational constraints within which the technology will operate [4, 7]. In addition, new technology designs may call for new policy changes or create new tensions between corporate and public policy. Leveraging HCI design methods (e.g., rapid prototyping with users) during policy design and implementation might help to create policies that are more responsive and robust to tech developments [11]. In sum, designing technologies and policies simultaneously should create better policies and better technologies.

So, how should HCI design technologies and policies simultaneously? Existing HCI research agendas and curricula do not offer a clear answer. Instead, the *research discourses* around tech and policy design are often held separately, speaking to different audiences, venues, and fields of expertise [10, 19]. As a result, policies at best take shape as an afterthought to technology design: “*Emergent*

technologies were designed; users adopted them; and when disputes emerged, aggrieved stakeholders turned to public agencies or the courts to clean up the mess” [10]. In most other cases—Consider Uber’s violation of taxi regulations when first inventing their ride-sharing service—technology design innovations broke the law first (often under the name of “*technology disruption*”) and still struggle to grapple with the policy and societal impacts a decade later [2, 12].

This workshop brings together technology and policy design communities to set up an HCI research and education agenda towards designing technologies and technology policies simultaneously. Notably, the need for such a tighter, more proactive design-policy integration is not new [10, 11]. However, the scale of emergent technologies and issues of HCI concerns today—from recommenders on social media to issues of sustainability and social justice—increasingly require tech designers to work at a societal-level scale in which policies operate (more on this in chapter 2). A community discussion and a shared research agenda on design-policy integration are more timely than ever.

In this context, this workshop will:

1. Identify the disciplines and areas of expertise needed for tighter, more proactive design-policy integration to succeed. We are particularly interested in identifying relevant areas of expertise that have not been regularly involved in HCI conferences and discourses;
2. Launch a community of researchers, educators, and practitioners interested in designing tech and policy simultaneously. We hope this group will collectively help build stronger, long-term connections between HCI, design, and policy communities;
3. Synthesize challenges encountered and lessons learned in prior attempts to design tech and policy simultaneously;
4. Develop a set of long-term and short-term research priorities for making advances on these challenges;
5. Develop a set of long-term and short-term education priorities for training a new type of HCI practitioner who has the knowledge and skills to design tech and policy;
6. Publish findings from this workshop (as an HCI conference publication and/or policy white paper), sharing outcomes with broader sets of HCI and policy researchers, educators, and practitioners.

2 RELATED WORK

While the call for a tighter, more proactive integration of technology and policy design processes is not new, the research and practitioner discourses around tech and policy design have largely happened far from one another [10, 11]. Here, we provide some examples of technologies and their HCI issues that might particularly benefit from simultaneous technology and policy designs. We then highlight some existing connections between technology and policy design work. When executing this workshop, we will work to include those who work on these technologies and technology issues across discipline boundaries, and to expand the connections in between.

2.1 Society-Level-Scale Technologies and Concerns that Need Design+Policy

We argue that the scale of many technologies and HCI issues today requires tech designers to work at a societal level, where public policy plays a big role in shaping people’s behaviors. Such areas might particularly benefit from simultaneous technology and policy designs. For example:

- *Computational systems of societal-level scale*, such as recommenders and content moderation systems on global, social media platforms; pervasive sensing and AI for health; and AI + Internet of Things, which includes smart cities, workplaces, domestic environments, and other environments. The scale and reach of such systems require designers to think beyond individual interactions, systems, or user populations, and to consider their macro consequences and policy implications [7, 21]. In building a healthier social media environment, what new solutions might emerge when prototyping new content moderation AIs and new platform policies simultaneously? To include commercial-level wearable devices (e.g., Fitbit) in clinical decision-making, how should the privacy protection laws, tech company policies, and wearable health AI designs coordinate with each other? In identifying a preferred future for smart transportation systems, how might one design autonomous vehicles’ behaviors differently if they consider [9]?
- *Foundation models and their app ecosystems*, such as the many applications built upon GPT-3 or Dall-E. Unlike traditional third-party app ecosystems (e.g., those accessible via Alexa), Foundation Models (FMs) are easily accessible to those without programming or AI skills and can power an unprecedented range of text- and image-generation applications out-of-the-box [3]. These characteristics can cause FMs’ unintended consequences (e.g., algorithmic bias) to spread at an unprecedented speed and scale. How can public and corporate policies and new FM application design tools complement each other in proactively preventing such unintended consequences from occurring?
- *Wicked problems at a societal scale*, such as sustainability, AI fairness, data governance, and privacy, persuasive design, and the attention economy, among others. These problems resist easy agreement on problem definition, desired outcomes, best approaches, and simple technological solutions. Addressing them requires both technological designs that can operate at scale and law and policy to affect economic markets, oversee organizational practices, or enforce responsibility and accountability [6, 18]. How can approaches in technology design and policies at different levels inform and coordinate with each other and address wicked problems such as social justice and AI ethics?

2.2 Connections Underlie Existing Technology and Policy Design Activities

A few shared commitments have already emerged in current technology and policy design practices. In some cases, HCI and policy designers have had similar ideas about how to approach these problems. This workshop aims to deepen and expand these implicit connections.

- *Defining preferred futures and framing wicked problems.* Both design and policy researchers have extensively studied wicked problems that resist clear objectives and simple technological solutions. Both fields have worked to understand, define, and discuss wicked problems. One recent project exemplifies how tech and public policy design might join forces in addressing wicked problems. Gilbert et al. [9] started by exploring different criteria for “good” transportation systems (e.g., enhancing social mobility, easy infrastructural maintenance such as road wear, desirable environmental impact). Based upon these criteria, they “reverse engineered” both design requirements for Autonomous Vehicles’ routing algorithms and needs for AV public policy.
- *Participation of diverse users and stakeholders, especially those from vulnerable populations.* In both technology design and policy, societal-scale problems pose questions about participation, power, and politics: Who benefits and who is harmed when addressing these problems? Who should be consulted or involved in looking for solutions? [6, 18] Both fields have turned to participatory design and co-design workshops as a way to include a richer set of stakeholders.
- *Anticipating possible unintended consequences.* Exploring how diverse users might use and misuse technologies are central to modern HCI design practices. What “reasonable actors” and “persons having ordinary skill” might expect from technology and how they might experience it as useful and valuable is central to policy-making and implementation work [10]. Can these activities inform and coordinate with each other in order to anticipate and address technologies’ unintended consequences?
- *Choreographing technology, people, and economic value flows.* Service designers routinely design how people and computational systems might collaboratively co-create value with customers and stakeholders. Service designers routinely ask who and when will need or not need to pay for a service. HCI scholars research the future of work and search for better divisions of responsibility between people and computer systems. We see an opportunity to bridge these practices and research with policy design work, especially those related to technologies’ economic policies, legal responsibilities, and labor protection laws.
- *Empirical research providing evidence of a need for technology and policy design.* Much current HCI research hopes to influence laws and regulations. Research on AI explainability and accountability offers one example area. However, this research often fails to account for what relevant government agencies and policymakers consider as acceptable evidence [19, 20]. The diverse methods and ways of knowing in HCI further complicate this research-to-policy translation [19]. Consider clinical research and policy as an analogy: The U.S. Food and Drug Administration codifies a set of evidence researchers need to provide in sequence in order to get a new drug approved eventually. Clinical researchers conduct drug research studies in this order, from animal studies to randomized controlled trials on real patients. Can HCI research and technology policy coordinate in analogous ways? What might a simultaneous HCI design and policy design process look like? Will it be more effective at generating safe and valuable new technologies?

3 ORGANIZERS

The co-organizers all have experience working at the intersection of technology and policy design. Our core areas of expertise include HCI and service design (Junginger, Yang, Zimmerman), technology public policy and ethics (Gilbert, Jackson, Wong), and legal design and design in policy making and policy implementation (Hagan, Junginger).

Thomas Krendl Gilbert is a Postdoctoral Fellow at Cornell Tech’s Digital Life Initiative. His research focuses on the emerging political economy of autonomous AI systems, including their technical development, moral significance, and implications for public policy [5]. His recent work investigates how specific machine learning procedures (e.g., reinforcement learning [8]) reframe classical ethical questions and recall the foundations of democratic political philosophy, namely the significance of popular sovereignty and dissent for resolving normative uncertainty and modeling human preferences. His work derives concrete implications for the design of AI systems that are fair for distinct subpopulations, safe when enmeshed with institutional practices, and accountable to public concerns, including medium-term applications like automated vehicles [9]. Dr. Gilbert served as the inaugural Law and Society Fellow at the Simons Institute for the Theory of Computing. He is also a research affiliate with the Center for Human-Compatible AI and co-founder of GEESE.

Margaret Hagan is the Executive Director of the *Legal Design Lab* and a lecturer at Stanford Law School and the Stanford Institute of Design (the d.school). She is a lawyer and holds a J.D. from Stanford Law School, a DPhil from Queen’s University Belfast, an M.A. from Central European University, and an A.B. from the University of Chicago. She specializes in the application of human-centered design to the legal system, including the development of new public interest technology, legal visuals, and policy design. Her research and teaching focus on the development and evaluation of new interventions to make the legal system more accessible. Professor Hagan has led workshops, including the “*Law + Design =*” summit at the Stanford Law School in 2017 that train legal professionals in the design process in order to produce client-focused legal design innovation. In addition, Professor Hagan teaches a series of project-based legal design classes, with interdisciplinary student groups tackling legal challenges through user-focused research and the design of new legal products and services.

Steven J. Jackson is an Associate Professor in the Department of Information Science and Department of Science and Technology Studies at Cornell University. He teaches and conducts research in the areas of scientific collaboration, technology ethics law and policy, democratic governance, and computational sustainability. More specifically, he studies how people organize, fight, and work together around collective projects of all sorts in which technology plays a central role. He also studies how infrastructure – social and material forms foundational to other kinds of human action – gets built, stabilized, and sometimes undone. This brings him regularly into worlds of policy, organizational or institutional analysis, and occasionally into the design. He spends much of his time doing ethnographic, legal, and sometimes historiographic research, where

he studies how shifting policies, emerging technologies, and cultural innovation meet complex and historically layered fields of practice. He thinks a lot about governance – how order is produced and maintained in complex sociotechnical systems; time – how we experience, organize, design, and work around the flows and patterns that shape and define individual and collective activity in the world; and breakdown, maintenance and repair, as crucial sites of inequality, power, and ethics in complex sociotechnical systems.

Sabine Junginger is a faculty at Lucerne School of Art and Design. Professor Junginger's research argues that taking policy-making and policy implementation as problems of design and as activities of design is critical for the success of policy innovation in the digital age [11]. She has explored bringing product and service design activities to policy innovation, such as identifying policy needs via future-oriented envisioning, iteratively prototyping policies, and evaluating outcomes. More broadly, Professor Junginger's core expertise is in the principles, methods, and processes of human-centered design, with a focus on design theories and design practices relevant to public and private organizations. At LUASA, Prof. Dr. Junginger heads the Competence Centre for Design and Management and co-leads the focus area *Organizations, HR and Leadership* for the LUASA-wide Interdisciplinary Theme Cluster (ITC) Digital Transformation of the Working World.

Richmond Wong is an Assistant Professor of Digital Media in the School of Literature, Media, and Communication at the Georgia Institute of Technology. His research seeks to understand how different forms of action can create value change and ethical outcomes in technology design, including new design practices, worker and community-led actions, organizational practices, and law and policy. Recent projects include studying the technology workers' practices related to ethics, and how companies have responded to new data protection laws. In addition to studying existing practices in these areas, he creates new activities and tools to help technology practitioners, users, and other stakeholders discuss and reflect on the social values and ethical issues related to technology.

Qian Yang is an Assistant Professor at Cornell University's Department of Information Science, with a graduate field appointment in Computer Science. Professor Yang is a human-AI interaction designer and researcher. Her research investigates how to bring human-centered thinking to bear on evermore complex AI systems. For example, her clinical decision support system design aided cardiology teams in better selecting artificial heart implant candidates. Her recent work focuses on designing societal-level-scale AI systems, such as GPT-3 and the natural language generation applications it powers, pervasive sensing devices and AIs as clinical decision support systems, and autonomous vehicles and their collective interactions with other road users. By designing such systems firsthand, Yang creates new design methods, processes, and tools that help HCI practitioners to better harness large-scale AI as design material while accounting for its unintended consequences.

John Zimmerman is the Tang Family Professor of AI and HCI at HCI Institute within Carnegie Mellon's School of Computer Science. He researches and designs human-AI interaction, human-robot interaction, and methods of innovating AI products and services. For

more than twenty years, Professor Zimmerman has designed novel, intelligent systems ranging from one of the first TV show recommenders to a crowd-sourced, transit arrival system to a decision support tool for implanting mechanical hearts to a system that keeps parents from forgetting to pick up their children. He has published more than 150 papers and is a member of the ACM CHI Academy. He teaches courses in service design, lean startup, and the design of AI products and services. While working for Philips, he invented the way everyone scrolls on their smartphone.

4 WORKSHOP PLANS

4.1 Pre-Workshop Plans

The workshop seeks to bring together a diverse group of researchers, educators, and practitioners who share our belief that now is the time for HCI to consider and act on policy as a core aspect of the work we do. Towards this goal, we will first advertise the workshop to the researchers, educators, and practitioners from relevant HCI and HCI-adjacent fields, such as design, law, public policy, policy design, public sector innovation, and political science. We, the co-organizers, come from these communities and will activate our professional networks (e.g., the AHRC Policy Network UK, the Design Research Society, the Danish Design Center, etc.) to draw in participants and invited speakers working on the intersection of policy, design, and technologies.

Before the Workshop, we will invite workshop candidates to submit position papers (1-2 pages) and short research papers (4-8 pages). Through short research papers, candidates can share their prior research and practical experiences in designing tech and policy simultaneously and reflect on lessons learned. In contrast, position papers can be think pieces where candidates share personal insights into the opportunities and open questions around marrying tech and policy design. Based on these submissions, we will select workshop participants with the aim of covering diverse areas of expertise, perspectives, and technology backgrounds to explore their policy design dimensions. We will share the selected workshop papers among all participants before the week of the workshop, providing a common ground for their discussion.

4.2 Remote/onsite Plans:

The workshop will be held with both in-person and remote options. The hybrid format will help attract participants and speakers who are new to CHI or may not be able to travel.

4.3 Workshop Structure

Our workshop activities are flexible in size. They can accommodate 16 to 50 participants depending on the size of the space and the number of in-person and remote participants. The workshop will employ 8 to 12-person working groups. The minimum number of working groups we need is two. We will create additional working groups if the number of participants exceeds 24. The organizers will attend the workshops, and we will always have at least one of the organizers playing the role of facilitator for each working group.

Session 1: Lightning Talks. Session 1 will have a brief welcome. This will immediately be followed by lightning presentations. Each participant will be given a few minutes to present the gist of

their position paper and detail why they have an interest in this workshop. The amount of time allocated for each presentation will be adjusted based on the number of participants who register. If the workshop is larger than 24, then we will run two parallel sessions. Presentations will be organized based on the themes drawn out of the submitted position papers.

The lightning talks will provide participants to get to know a little bit about one another. This should build a little trust among participants, making later workshop activities more effective. In addition, it will provide a first view of everyone's interests which will be important during the final session.

Coffee Break

Session 2: Brainstorming Round 1. For session 2, we will break participants into working groups based both on the thematic fit of position papers and on their interest in working on HCI education priorities or HCI research priorities. A member of the organizing committee will facilitate a brainstorming session to draw out and discuss different priorities. The collected set of ideas will be ranked in terms of their priority to the participants. In addition, they will be classified as long-term or short-term, and they will detail dependencies among the different ideas.

The working groups will reform into larger groups. Each will share the priorities they identified and their rationale for rankings and classifications. We will have a short discussion to see if there is collective agreement and to discover new perspectives on the situation.

Lunch

Session 3: Brainstorming Round 2. For session 3, we will form new working groups, allowing participants to self-select the theme and topic (education priorities or research priorities) they want to work on. We will engage in a second round of brainstorming facilitated by workshop organizers. Synthesis of the ideas will deviate from the first round in that we will push participants to provide more detailed rationales for the rankings and the classification. We expect this round will reproduce many of the ideas from round 1 as well as new ideas that come from some personal reflection and from the larger group discussion. Once again, the working groups will come back together and share and discuss their sets of priorities.

Coffee Break

Session 4: Collaboration Speed Dating. During the coffee break, organizers will collect lists of matches between participants. These are explicit indications of other participants a participant wishes to talk to about future collaboration. Based on their indications, we will generate a short list of one-on-one collaboration meetings.

During the speed dating session, participants will have 8 minutes to sit together and talk about future collaboration. At the end of the eight minutes, they will rotate to the next name on their speed dating list. Our hope for this activity is that at the end of the workshop we will have facilitated the genesis of several new collaborations around the topic of policy and HCI.

Session 5: Wrap-up. In the final few minutes of the workshop, we will hold a group discussion about writing a paper describing the workshop and the priorities that emerged. We will collectively discuss the best venue for the paper and collect the names of everyone who wishes to participate in writing. We were particularly

impressed with the impact of a paper on social computing that emerged from a workshop and had participation from almost all who attended [13]. We hope to have a similar impact from this post-workshop activity.

4.4 Post-Workshop Plans

We will work to share the workshop outcomes with broader sets of HCI and policy researchers, educators, and practitioners. Specifically, we will publish an HCI conference publication (e.g., at the Communications of the ACM) and/or a policy white paper, describing (1) a synthesized set of challenges that emerged and lessons learned in prior attempts to design tech and policy simultaneously; (2) a set of long-term and short-term research priorities for making advances on these challenges; and (3) a set of long-term and short-term education priorities for training a new type of HCI practitioner who has the knowledge and skills to design tech and policy.

In addition, we will explore additional venues to continue building and sustaining long-term connections between HCI and adjacent policy design communities in accounting for technologies' unintended consequences.

5 CALL FOR PARTICIPATION

Accounting for technologies' unintended consequences—whether they are misinformation on social media or issues of sustainability and social justice—increasingly requires HCI to consider technology design at a societal-level scale. At this scale, public and corporate policies play a critical role in shaping technologies and user behaviors. How can technology design and policies better inform and coordinate with each other in generating safe new technologies? What new solutions might emerge when HCI practitioners design technology and its policies simultaneously to account for its societal impacts? This one-day, hybrid workshop brings together HCI, design, and policy communities to explore these questions.

Submissions may address one or both questions in a position paper (1-2 pages) or a short research paper (4-6 pages). Submissions may focus on broad insights into the opportunities and open questions around designing tech and policy simultaneously. Submissions can also offer specific research or case studies. We especially encourage submissions related to society-level-scale technologies and HCI concerns where design+policy can be particularly valuable. Submissions should use the single-column ACM template.

The submission link will be available on the [workshop website](#): designpolicy.one. We also plan to make the accepted workshop papers available on the website. All submissions will be reviewed by the workshop's co-organizers and other expert invited reviewers. Authors of the highest-rated submissions will be invited to give lightning talks (Session 1) with our invited speakers. At least one author of each accepted submission must attend the workshop. All participants must register for both the workshop and for at least one day of the conference.

ACKNOWLEDGMENTS

Qian Yang's effort to help organize and participate in the workshop was partially supported by National Science Foundation under grants IIS-2107111 and IIS-2212431. John Zimmerman's effort to

help organize and participate in the workshop was partially supported by National Science Foundation under grants IIS-2007501 and IIS-2112633. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Sabine Junginger's work on integrating policy and service design is supported by the Swiss National Science Foundation SINERGIA CRSIIS 189955 VA-PEPR (Voice Assistants, People Experiences, Practices and Routines).

Ideas of this workshop were partially developed during the 'AI and Humanity' 2022 summer program at the Simons Institute for the Theory of Computing. Two co-organizers, Qian Yang and Thomas Gilbert, were supported by the Simons Institute to participate in the summer cluster.

REFERENCES

- [1] 2010 Americans with Disabilities Act (ADA) Standards for Accessible Design. 2014. U.S. Department of Labor. https://www.ada.gov/2010ADAsstandards_index.htm
- [2] Justin Bariso. 2016. Why does uber keep breaking the law? because they're disrupting, of course. <https://www.inc.com/justin-bariso/why-does-uber-keep-breaking-the-law-because-theyre-disrupting-of-course.html>
- [3] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khaftab, Pang Wei Kohd, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray O gut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. 2021. On the Opportunities and Risks of Foundation Models. [arXiv:2108.07258 \[cs.LG\]](https://arxiv.org/abs/2108.07258)
- [4] Alissa Centivany. 2016. Policy as Embedded Generativity: A Case Study of the Emergence and Evolution of HathiTrust. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing* (San Francisco, California, USA) (CSCW '16). Association for Computing Machinery, New York, NY, USA, 926–940. <https://doi.org/10.1145/2818048.2820069>
- [5] Roel Dobbe, Thomas Krendl Gilbert, and Yonatan Dov Mintz. 2020. Hard Choices in Artificial Intelligence: Addressing Normative Uncertainty through Sociotechnical Commitments. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* (2020).
- [6] Lynn Dombrowski, Ellie Harmon, and Sarah Fox. 2016. Social Justice-Oriented Interaction Design. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*. ACM Press, New York, New York, USA, 656–671. <https://doi.org/10.1145/2901790.2901861>
- [7] Brianna Dym, Namita Pasupuleti, and Casey Fiesler. 2022. Building a Pillowfort: Political Tensions in Platform Design and Policy. *Proc. ACM Hum.-Comput. Interact.* 6, GROUP, Article 16 (jan 2022), 23 pages. <https://doi.org/10.1145/3492835>
- [8] Thomas Krendl Gilbert, Sarah Dean, Tom Zick, and Nathan Lambert. 2022. Choices, Risks, and Reward Reports: Charting Public Policy for Reinforcement Learning Systems. (2 2022). <https://doi.org/10.48550/arxiv.2202.05716>
- [9] Thomas Krendl Gilbert, Aaron J Snoswell, Michael Dennis, Rowan Mcallister, and Cathy Wu. 2022. Sociotechnical Specification for the Broader Impacts of Autonomous Vehicles. *ICRA'22 workshop Fresh Perspectives on the Future of Autonomous Driving*.
- [10] Steven J. Jackson, Tarleton Gillespie, and Sandy Payette. 2014. The Policy Knot: Re-Integrating Policy, Practice and Design in CSCW Studies of Social Computing. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (Baltimore, Maryland, USA) (CSCW '14)*. Association for Computing Machinery, New York, NY, USA, 588–602. <https://doi.org/10.1145/2531602.2531674>
- [11] Sabine Junginger. 2013. Design and Innovation in the Public Sector: Matters of Design in Policy-Making and Policy Implementation. *Annual Review of Policy Design* 1 (2013). Issue 1. <https://ojs.unbc.ca/index.php/design/article/view/542>
- [12] Sarah Kessler. 2013. Uber: When innovation outpaces the law. <https://www.fastcompany.com/3001169/uber-when-innovation-outpaces-law>
- [13] Aniket Kittur, Jeffrey V Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. 2013. The future of crowd work. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 1301–1318.
- [14] Jonathan Lazar. 2010. Interacting with Public Policy. *Interactions* 17, 1 (jan 2010), 40–43. <https://doi.org/10.1145/1649475.1649485>
- [15] Jonathan Lazar. 2015. Public Policy and HCI: Making an Impact in the Future. *Interactions* 22, 5 (aug 2015), 69–71. <https://doi.org/10.1145/2807916>
- [16] Lawrence Lessig. 1998. WHAT THINGS REGULATE SPEECH: CDA 2.0 VS. FILTERING. *Jurimetrics* 38, 4 (1998), 629–670. <http://www.jstor.org/stable/29762575>
- [17] Kim Lyons. 2022. Twitter rolls out its alt badge and improved image descriptions. <https://www.theverge.com/2022/4/7/23015669/twitter-alt-badge-improved-image-descriptions-accessibility>
- [18] Robert Soden, Pradnya Pathak, and Olivia Doggett. 2021. What We Speculate About When We Speculate About Sustainable HCI. In *ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS)*. ACM, 188–198. <https://doi.org/10.1145/3460112.3471956>
- [19] Anne Spaar, Abigail Durrant, Chris Elsden, and John Vines. 2019. *Understanding the Boundaries between Policymaking and HCI*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3290605.3300314>
- [20] Anne Spaar, Nick Spencer, Abigail Durrant, and John Vines. 2022. Creative and collaborative reflective thinking to support policy deliberation and decision making. *Evidence & Policy: A Journal of Research, Debate and Practice* 18 (04 2022). <https://doi.org/10.1332/174426421X16474564583952>
- [21] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-Examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376301>