



The Future of HCI-Policy Collaboration

Qian Yang
Cornell University
USA
qianyang@cornell.edu

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

Steven J. Jackson
Cornell University
USA
sjj54@cornell.edu

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

Margaret Hagan
Stanford University
USA
mdhagan@stanford.edu

Thomas Krendl Gilbert
New York Academy of
Sciences
USA
tgilbert@nyas.org

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

ABSTRACT

Policies significantly shape computation’s societal impact, a crucial HCI concern. However, challenges persist when HCI professionals attempt to integrate policy into their work or affect policy outcomes. Prior research considered these challenges at the “border” of HCI and policy. This paper asks: What if HCI considers policy integral to its intellectual concerns, placing system-people-policy interaction not at the border but nearer the center of HCI research, practice, and education? What if HCI fosters a mosaic of methods and knowledge contributions that blend system, human, and policy expertise in various ways, just like HCI has done with blending system and human expertise? We present this re-imagined HCI-policy relationship as a provocation and highlight its usefulness: It spotlights previously overlooked system-people-policy interaction work in HCI. It unveils new opportunities for HCI’s futuring, empirical, and design projects. It allows HCI to coordinate its diverse policy engagements, enhancing its collective impact on policy outcomes.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; *HCI theory, concepts and models*; • **Computing methodologies** → *Artificial intelligence*.

KEYWORDS

Policy, design, societal impact of technology.

ACM Reference Format:

Qian Yang, Richmond Y. Wong, Steven J. Jackson, Sabine Junginger, Margaret Hagan, Thomas Krendl Gilbert, and John Zimmerman. 2024. The Future of HCI-Policy Collaboration. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3613904.3642771>

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.

CHI '24, May 11–16, 2024, Honolulu, HI, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0330-0/24/05

<https://doi.org/10.1145/3613904.3642771>

1 INTRODUCTION

From the future of work, to equity and sustainability, computation’s societal impact has become a critical HCI concern. At least three forces help shape this impact: the incentives, deterrents, and rules wired into computational systems (“system design”); the social and cultural tendencies of people (“social practice”); and the laws and regulations that govern both systems and people (“policy”) [20, 43, 47, 96]. Take privacy as an example: The design of a mobile app, its users’ privacy awareness, and related policies (e.g., app stores’ mandates, governments’ regulatory fines) all help decide how well this app will protect user privacy [42].

In many cases, HCI researchers and practitioners are already successful in improving system-people-policy interaction. For instance, HCI scholar Lorrie Cranor both conducted empirical and design research on privacy, and led the drafting of privacy regulations at the U.S. Federal Trade Commission (FTC) [24]. Professor Colin Gray’s work on deceptive web design patterns influenced the U.S. DETOUR (Deceptive Experiences To Online Users Reduction) Act [5, 36]. Such success stories span many countries and social issues, such as content moderation [39], gig worker protection [61, 67, 85], AI safety [87], smart cities [26, 91], sustainability [14], and more.

Yet in many other cases, policy considerations lag behind, or remain isolated from) the other HCI work that centers around systems and people. Consider Uber’s creation of gig work (and intentional breach of taxi regulations) when it first launched its ride-sharing service [12, 51], Facebook’s adoption of AI newsfeed rankers (and amplification of misinformation) [39], or OpenAI’s public release of chatGPT (and intentional violation of copyright laws) [83, 92]. Too often, HCI seemed to become excited about the technology first, take an interest in policy only after things had gone off the rails, and then find itself in the unenviable position of trying to put genies back into bottles.

This paper aims to better understand this disparity and amplify HCI’s collective voice in the policy realm. From difficulties in bringing policymakers to the table [77], to misalignment between HCI methods and policymakers’ evidentiary needs [78], many challenges of HCI-policy collaboration are well-known. Prior work referred to these challenges as at the “boundaries” of HCI and policy [77], prompting decade-long calls for more collaborations across HCI-policy disciplinary boundaries [20, 24, 43, 47, 57, 58, 86]. These calls are necessary and valuable [60]; nothing in this paper argues

against them. Meanwhile, the decade-long persistence of these calls also raises the question: Can fresh perspectives on these challenges inspire additional strategies for overcoming them? This paper aims to provide one such perspective.

This paper argues that—thanks to the decade-long calls for more HCI-policy collaboration—HCI has begun creating new methods and new types of knowledge contributions that are distinct from traditional HCI or policy ones. HCI and policy expertise have begun cross-pollinating. However, such unconventional, trans-disciplinary work has seemed to struggle with getting past reviewers, therefore has remained largely unrecognized in peer-reviewed HCI research venues, industry practices, or HCI textbooks. This lack of recognition can contribute to the disparity in HCI’s policy engagement and impact.

Reframing the challenges of HCI-policy collaboration as “*insufficient recognition of trans-disciplinary methods and knowledge*” informs new strategies for addressing them: What if HCI were to see policy as integral to its intellectual concerns, placing system-people-policy interactions not “*at the boundaries of HCI and policy*” [77], but nearer the *center* of HCI research, practice, and education? What if HCI were to foster a mosaic of distinctively HCI methods and knowledge contributions that blended system, human, and policy expertise in varying ways and degrees, just like it has previously done with blending system and human expertise?

This paper sketches out this re-imagined HCI-policy landscape as a *provocation*. We demonstrate how it can (1) spotlight previously overlooked system-people-policy interaction work and knowledge, (2) unveil new opportunities for HCI’s futuring, empirical, and design projects, and (3) allow HCI to coordinate its diverse policy engagements, thereby enhancing HCI’s collective impact on policy outcomes. We show how it opens up many new research questions, therefore is a useful seed for broader, HCI-community-wide discussions. We invite fellow HCI researchers and practitioners to discuss their views of the relationship among systems, people, and policy expertise, and how HCI wants to position itself in this relationship.

This paper makes three contributions. First, it moves the research discourse on HCI-policy collaboration beyond the decade-long *calls for more collaboration*, towards a community-wide discussion on its problem-solution framing. Second, it jump-starts this discussion by proposing one new problem-solution frame: There is insufficient recognition that policy is not just HCI’s “*broader impact*”, but integral to HCI’s intellectual pursuits in system-people interaction. Finally, it exposes important new questions for future research to address. For example, when is policy *not* the best solution to tech’s societal harms? And when sudden momentum for policy change appears, how might HCI quickly assemble its diverse work and seize the opportunity?

2 KEY CONCEPTS

At the outset of this paper, we plunged head first into the topics of “*policy*” and “*HCI-policy collaboration*” without detailing what they meant. These are nuanced concepts, variably interpreted and lacking agreed-upon definitions [49, 68, 89]. With these complexities in mind, we unpack what we mean by these terms. Our goal is not to establish conclusive definitions, but to scope them for the narrow purpose of this paper.

2.1 Policy

In this paper, policy refers to principles, guidelines, and written rules formulated and adopted by an authority, organization, or government [43]. Among the widely-accepted policy definitions, we chose this one because it has both the specificity and breadth that match our research goals.

First, it encompasses various rules external to a computing system that wield authority over human-computer interactions, such as governmental policies, tech platform policies, and more. All these policies help shape computation’s societal impact, therefore hold relevance to HCI.

Secondly, this definition recognizes that a policy encompasses not only its written rules, but also the ways in which these rules are formulated and applied. Take the EU AI Act as an example. Whose inputs shaped this legal framework [69], how government agencies translate it into enforceable rules, and how the courts interpret and apply the rules to specific situations [90] all influence AI’s societal impact. All are relevant to HCI.

This definition of policy (and the scope of this paper) is also deliberately narrower than some prior work. *It does not intend to address all aspects of policy practices, nor all political forces that can influence computation’s societal impact*. Because our goal is to understand the challenges of HCI-policy collaboration, we focus on the aspects of policy uncommon in current HCI work. For example, a society’s cultural politics (e.g., its conception of fairness [63]) and governmental policy-making procedures both significantly influence its tech policies. Yet because the former is already common in HCI work, the latter takes precedence in this paper.

2.2 “*Understanding*” Policy, “*Designing*” Policy

What this paper refers to as *understanding* and *designing* policy reflects this policy definition. Because “*policy*” encompasses both the written rules and the process of their formulation and enactment, “*understanding*” a policy includes understanding both aspects. Similarly, “*designing*” policy includes not only formulating rules, but also designing the processes for formulating and implementing the rules [71]. In practice, policy design is the iterative process of (1) identifying policy needs, (2) clarifying policy needs (or issue-framing), (3) formulating policy, (4) designing systems and services that implement policy, and (5) evaluating policy outcomes [49, 70]¹. We chose a broader view of designing policy because we want to consider various HCI-policy connections [28, 94].

The synergies between technology and policy design are self-evident. Both processes start with recognizing an opportunity for new systems/policies to help, and proceed with prototyping and evaluating their helpfulness iteratively. These synergies have prompted some HCI researchers to call for designing technology and policy simultaneously [96] and some legal scholars to call HCI a new direction in designing Information Technology law [86].

¹This scope of policy “*design*” is intentionally broader than some prior HCI literature, which focused on evidence-based policy-making (and HCI as a supplier of evidence) [77]. It is also broader than some legal scholarship, which sees computing system “*design*” as merely the implementation of a given policy requirement (e.g., generating a standard cookie opt-out menus [84], rather than a process of open-ended knowledge inquiry.

2.3 “Productive” HCI-Policy Collaboration

By definition, HCI addresses the interactions between computing systems and people. As these systems grew ubiquitous and integral to everyday life, HCI has expanded its focus from single-user-single-system interactions, to interactions among platforms, stakeholders, organizations, societal processes, and even planetary concerns [14, 32, 95]. Policy entered HCI discourse during this expansion, because it significantly shapes these larger-scale and often political interactions [27, 43, 60].

A productive HCI-policy collaboration can improve computation’s societal impact, by catalyzing synergistic designs of computational systems, social practices, and policies [29, 43]. Consider, for example, a novel AI system that helps diagnose a disease (system), clinician teams’ diagnostic workflow (social practice), and clinical malpractice laws (policy). To improve diagnostic accuracy in practice, designers cannot simplistically match clinician workflows and malpractice laws with this new AI system, nor vice versa. Instead, the three designs need to co-evolve [70]. Computation’s societal impact improves step-wise, as the “knot” of system design, social practice design, and policy design loosens and tightens, unwinds, and reties [43].

There is no single best approach to HCI-policy collaboration that prescribes this ideal outcome [30]. Should HCI professionals generate empirical evidence of technological harm, before persuading policymakers to act [37]? Or should they envision a computational system and its regulations simultaneously [96]? Should they seek a government position or work with advocacy groups when prompting new policies? Depending on context, the most fruitful approach to HCI-policy collaboration varies, each bringing different challenges.

2.4 Problem-Solution Frames

How might we grasp the challenges of HCI-policy collaboration and devise solutions? This is a wicked problem [73]. Because challenges facing every collaboration vary, there is no definitive answer to the question of “*which challenges(s) are responsible for the disparity in their success levels overall*” or “*what solutions might help*” [27].

Addressing wicked problems relies on problem framing and reframing [73]. By seeing a problematic situation through different lenses, one can better understand its various dimensions and see new avenues for potential solutions [28]. And *this* is the goal of this paper. We wanted to find a new frame for understanding the challenges and disparity in HCI-policy collaboration; a frame that can reveal new avenues for solutions.

3 BACKGROUND

HCI-policy collaboration has many well-known success stories and persistent challenges. For example: difficulties in bringing policymakers to the table [77], policymakers demanding policy evidence that (1) may not exist (e.g., evidence of a new technology’s not-yet-manifested societal form) or (2) may not be the types of evidence HCI approaches produce [77, 78]; mismatch between the pace of policy or political changes and the pace of HCI work [24, 66]; and more.

Interestingly, prior research rarely deliberated on the nature of these problematic situations. Instead, it directly suggested solutions.

Most commonly, researchers called for more HCI-policy collaboration [20, 24, 43, 47, 57, 58, 86]. Others suggested more concrete solutions. For instance, to bring policymakers to the table, they encouraged more HCI scholars to take a gap year and work at government agencies [59]. To address the HCI-policy evidentiary gap, they recommended involving policymakers in participatory design workshops, to help them appreciate HCI methods [77]. To align the misaligned timelines of policy and HCI work, HCI researchers recommended committing to policy work for the long run and adjusting to the inherently slower pace of policy change [24]. These recommendations yielded valuable results [60]; nothing in this paper argues against them. Nevertheless, we see opportunities to deliberate on the problem-solution frames that these solutions imply. Let us illustrate these opportunities through two examples.

3.1 Current Problem-Solution Frame #1

The numerous calls for more HCI-policy collaboration suggest that a critical hindrance to this collaboration is that too few HCI professionals are participating [57].

However, the situation might be shifting. Over the past decade, methods such as participatory design and value-sensitive design—methods many policy actors also use—have moved more towards the center of HCI [33, 74]. Since 2018, the ACM FAccT conference has been bringing together HCI, law and policy, and other fields to address AI’s ethical issues [56]. Since 2021, SIGCHI publications mentioning “policy” surged by more than 40%, according to ACM Digital Library. In this context, it is worth asking whether HCI’s lack of attempts to policy engagement remains true today.

3.2 Current Problem-Solution Frame #2

HCI researchers have characterized the challenges of HCI-policy collaboration as occurring at the boundaries of HCI and policy [77]. This framing sees policy change as a “*broader impact*” of HCI’s intellectual work (i.e., understanding and designing system-people interaction). Therefore, the disciplinary boundaries between HCI and policy can hinder their collaboration.

Along this line, prior research repeatedly recommended that HCI and policy communities accept each other’s norms. For example, methodologically, HCI researchers recommended that policymakers embrace HCI’s design methods [77]; socially, they encouraged more HCI professionals to socialize with policy actors [57, 59]; temporally, they encouraged HCI researchers to adapt to policy and political timelines [24].

These suggestions are highly valuable, but seem to be partial solutions for fostering productive HCI-policy partnerships. Take, for instance, the timeline of addressing the societal impact of generative AI (genAI). How should HCI and policy communities coordinate their pace of work to best address genAI’s societal impact? How should HCI communities approach publishing novel GPT applications, considering some *might* be soon regulated out of existence? These questions require nuanced debates [41]. Urging one community to adopt the other’s timeline appears insufficient.

More fundamentally, a productive collaboration between any two disciplines involves more than choosing whose norms to conform to. It entails creating new *trans-disciplinary* methods and new bodies of knowledge distinct from either parent discipline [44, 54]. The

innovations of computing systems, social practices, and policies require distinct methods, and each encompasses multiple temporal patterns [64, 80, 81]. Are there trans-disciplinary methods that can respect and bridge these differences? Such meta-discussions are absent in current HCI-policy research.

4 RESEARCH ACTIVITIES

We (a group of HCI, design, law, and policy researchers) wanted to reframe the challenges facing HCI-policy collaboration, in the hope of identifying new opportunities for addressing them. Towards this goal, we collected empirical reports of previous HCI-policy collaboration processes, analyzed factors contributing to their varying levels of success, and then worked to identify a new, useful problem-solution frame.

4.1 Literature Review

We first searched peer-reviewed HCI research publications and practitioner-facing books for empirical reports of previous and ongoing collaboration processes. However, we soon realized that this literature has published very few such reports. Even the most fruitful collaborations—collaborations that resulted in national policies with HCI’s scholars’ names [24, 36]—left little documentation of how HCI researchers and practitioners worked with policy actors in practice, or how they approached the gap between HCI and policy.

HCI’s Participatory Design (PD) projects are telling examples. Peer-reviewed HCI venues have published many such projects, where researchers investigated vulnerable stakeholders’ needs and engaged them in drafting policy recommendations [91, 97]. While these publications offered detailed need-finding and rule-making processes, it is difficult to gauge whether or how they influenced policy outcomes. After all, citing academic references or attributing individual researchers is uncommon in many legal or public policy documents.

A larger set of HCI-policy collaboration efforts reside outside of peer-reviewed HCI publications. We found their traces in less prominent genres of HCI research dissemination, e.g., in <Interactions> magazines [58, 66], in non-archival policy white papers [83], in the news [18], and on Medium [50]. Unfortunately, these publications also provided little detail about the HCI-policy collaboration contexts, processes, or outcomes.

Other collaboration efforts seemed entirely untraceable. For example, we suspect HCI research has influenced Facebook’s content moderation policies, because many HCI scholars worked on the topic as Facebook employees. However, such collaboration is difficult to verify, much less to study rigorously.

4.2 Community Inputs and Discussions

To collect more data on HCI-policy collaboration processes and breakdowns, we held a workshop on this topic at the 2023 CHI conference. With an open call for participation, the workshop brought together 57 HCI researchers and practitioners from the Americas, Asia, Australia, and Europe. The workshop received and published more than 60 papers detailing the participants’ specific research projects, all at some intersection of computational systems, people, and policy. They cover an overwhelming breath of topics, e.g., social media’s role in human trafficking, surveillance of migrant

workers, unions negotiating tech policy for workers, combined use of HCI and policy in tenant protection, policy issues throughout the supply chain of generative AI, AI in e-governments, challenges in applying GDPR to user interface design, and many more. These publications became an additional source of data for our analysis.

4.3 Data Limitations

One limitation of our data is geographical. While workshop participants represent many regions globally, all authors of this paper are based in the U.S. and Europe, as are most workshop publication authors. All except one workshop author are based in democracies. Recognizing this geographical bias, we strongly encourage researchers from other parts of the world to help critique this work and share their perspectives on HCI-policy collaboration.

4.4 Data Synthesis

We synthesized a new problem-solution frame based on this relatively diverse set of empirical reports. As we will demonstrate in the remainder of the paper, this frame can (1) help explain the observed disparity in HCI-policy collaboration and (2) reveal new avenues of opportunity in addressing the disparity. It is (3) flexible, allowing HCI researchers and practitioners to derive specific actions according to the respective context of their policy engagement. Finally, even the authors lack consensus on whether this problem-solution frame is too progressive and controversial, or too obvious that is what HCI needs to do. In this sense, this frame can be (4) an effective seed for broader HCI-community-wide discussion.

This frame results from the authors’ year-long discussion within our small team, with inputs from wider our respective communities. Appendix A describes our deliberation process in more detail.

5 POLICY AS INTEGRAL TO HCI’S INTELLECTUAL CONCERNS

5.1 A Vision of the Future

We envision a future in which HCI sees policy as integral to its intellectual pursuits, placing system-people-policy interaction not at the boundaries of HCI and policy, but nearer the center of HCI research, practice, and education. This vision differs from prior framings in three important ways.

1. HCI will deliberately dissolve the boundary between its policy- and system-human-interaction-focused work. Instead, a wealth of hybrid methods and hybrid knowledge contributions will bridge and harmonize HCI’s system, people, and policy expertise.
2. Each HCI project will choose to integrate system, people, and policy considerations to varying degrees and in diverse modes. For example, some projects will advance HCI’s knowledge of system-human interaction, while providing policy implications. Some projects will leverage human-system interaction expertise and advance HCI’s knowledge of policy design. Other projects will generate new knowledge on system-human-policy interaction, by bridging the three areas of expertise. In this future, peer-reviewed HCI research venues will critically assess and accept these diverse knowledge contributions. HCI education and practitioner methods will embrace these varied methods.

3. HCI communities will actively coordinate these varied methods and diverse knowledge contributions, maximizing HCI's collective impact on real-world policy outcomes.

5.2 A New Problem Frame

There are three key differences between this vision of the future and the current landscape of HCI-policy collaboration. These differences offer us new insights into why HCI-policy collaboration efforts have had disparate results.

1) Lack of recognition for the inherent tensions between HCI and policy methods, leading to a lack of recognition for trans-disciplinary methods that address the tensions. Peer-reviewed HCI research venues, practitioner-facing methods, and HCI textbooks did not always give due recognition to unconventional methods and knowledge that bridged HCI-policy expertise. For example, while many peer-reviewed CHI and CSCW papers feature policy recommendations derived from PD workshops (a well-established HCI method) [33, 82], none featured policy proposals derived from researchers' own synthesis (for example, by synthesizing policy-makers' evidentiary needs and HCI's prior empirical work) [24]. This may be because HCI has not yet acknowledged synthesis as an established HCI research method, or policy design as an HCI knowledge contribution. Consequently, methods that can bridge HCI-policy evidentiary gap remain largely visible in prominent HCI venues.

This lack of recognition can contribute to the disparity in HCI-policy collaboration. It can disincentivize early-career researchers from participating. It also makes HCI-policy knowledge less accessible, adding to why socializing with policymakers *appears* to be the only way to gain such knowledge [24, 59, 77]. Both effects privilege established HCI experts (who are more likely to prioritize long-term policy impact and relationship-building over near-future publications). Both privilege issues aligned with policymakers' existing interests over precautionary debates.

2) Challenges in deliberating and curating novel HCI-policy methods and knowledge contributions. The limited visibility of novel HCI-policy methods and contributions not only hinders methodological innovation. It also inhibits broader HCI communities from scrutinizing or deliberating upon these emergent methods and findings. For example, while many Speculative Design and PD projects have discussed their implications for policy, few have asked: Are "implications for policy" a necessary or good measure of the quality of empirical work that studies technological harm? What makes an implication for policy more meaningful or useful than others?

The lack of deliberation on emerging HCI-policy methods means fewer proven method choices for HCI professionals attempting to integrate policy into HCI work. Little publicly available guidance exist on how to write better "Implications for Policy." HCI practitioners' toolbox contains few methods for designing system-people-policy interactions. This lack of public knowledge yet again disadvantages early career HCI professionals.

3) Insufficient community-wide coordination for collective impact. Efforts to coordinate HCI's diverse modes of policy engagement are scarce in HCI literature. Considering that policies underpin every computational system and every human, when

is policy design *not* the best approach to improve a system's human impact? Knowing that momentum for policy change can arise abruptly, how can HCI quickly mobilize its diverse HCI-policy work to seize the opportunity? Answers to these meta-questions can amplify HCI's collective influence on policy outcomes, yet are noticeably absent in today's HCI research discourse.

With little community-wide coordination and support, HCI researchers and practitioners relied on their respective efforts for policy impact, exacerbating disparities in HCI-policy collaboration and weakening the collective impact of HCI on policy.

5.3 A New Solution Frame

These new problem frames reveal new solution frames; new strategies for addressing the disparity in HCI-policy collaborations.

1. Explicate underlying tensions between HCI and policy methods;
2. Give due recognition to the trans-disciplinary methods that effectively address these tensions. Assess and accept their knowledge contributions to peer-reviewed HCI research venues, practitioners' toolboxes, and HCI textbooks;
3. Foster a mosaic of trans-disciplinary methods and knowledge contributions that blend policy and HCI's futuring, empirical, and design expertise to varying degrees and in diverse ways. Each individual HCI project can choose among them;
4. Coordinate HCI's diverse policy engagements and maximize HCI's collective impact on policy outcomes.

The remainder of the paper demonstrates the usefulness of this new problem-solution frame. We illustrate how even seeing existing literature through this new frame can reveal tangible, new opportunities for individual HCI projects (Chapter §6) and for HCI communities as a whole (Chapter §7).

Do HCI communities endorse this re-imagined HCI-policy relationship? Are these strategies indeed effective in fostering HCI-policy collaboration and improving computation's societal impact? Addressing these questions takes time and requires community-wide HCI efforts and debates. The opportunities outlined initiate such efforts and debates.

6 OPPORTUNITIES FOR INDIVIDUAL HCI PROJECTS

Seeing policy as integral to HCI's intellectual endeavors—be it revealing the impact of existing technologies ("*empirical work*"), improving this impact by creating new technologies ("*design work*"), or speculating the societal impact of emerging technologies ("*futuring work*")—can enhance these endeavors.

6.1 Integrating Policy into Empirical Work

HCI-Policy Synergies and Tensions. Seeing policy as integral to HCI's pursuit in human understanding helps us see the connections and tensions between them. Both policy and HCI communities want to understand people's values, behaviors, societal processes, and interactions with and experiences of emergent technologies. HCI and policy share empirical methods such as participatory workshops [91, 97]. In many cases, they are already collaborating, e.g., in understanding misinformation on social media [39] and promoting gig worker welfare [61, 76].

Nevertheless, tensions may arise between HCI and policy actors' empirical work, because of their different scopes. HCI's empirical findings illustrate *situated* interactions among a specific combination of systems, people, policies, and contexts. They mean to inform system designs that meet the needs of specific stakeholders in particular contexts. In contrast, policy design influences a broader set of technologies, people, and interactions, thus requiring larger-scale empirical evidence. This difference in scope has led to some frustrations in the HCI community that policymakers seek difficult or impossible empirical evidence, such as large-scale evidence of an emergent technology's societal harm [77].

Low-Hanging Fruits in Addressing HCI-Policy Tensions. By explicating the tensions between HCI and policy's human understanding work, we can identify first steps in alleviating these tensions. For example, HCI empirical work can communicate the generalizability of its policy implications more explicitly. Is an observed technological harm specific to this particular technology or population? Is it the concern of global, national, local, or sectoral regulations? Does it require an update to the spirit, the text, or the implementation of the law? Answers to these questions are within the reach of existing HCI empirical methods, yet can address policymakers' needs more directly.

HCI can also expand its empirical methods, observing how a specific policy helps shape various human-computer interactions, thereby increasing the likelihood of deriving generalizable policy implications. What limitations of current AI regulations does generative AI expose? Is a new privacy policy in conflict with the social norms of particular populations, such that they use technologies to circumvent it?

Expanding HCI's use of metaphors to communicate policy implications represents another tangible opportunity. HCI communities have long used metaphors when communicating a technology's affordances and human impact. Less discussed is that metaphors often emphasize different social values and carry rich policy implications (e.g., Are generative AI models more like Internet search engines (therefore their data contributors enjoy similar rights) or black boxes? Is attaching a GPS surveillance device to a car more like following a car on public roads (therefore does not require a court-ordered warrant) or more like trespassing on one's private property? [2]). In tech policy discussions, various interest groups often debate the choice of metaphors, because this choice anchors how policymakers understand a technology's affordances and decide whether it needs new policy and legal frameworks [93]. There is a ready opportunity for HCI to improve its communication with policymakers via metaphors. It can help link HCI's contextual empirical insights to broader policy contexts, using a language that they are familiar with.

HCI's empirical work might also take advantage of the fact that law and policy differ across geographic scales and jurisdictions, linking local observations with broader policy implications. Take, for instance, the data breach notification laws in the U.S. Rather than a single national law governing data breaches, this set of patchwork rules was created and passed on a state-by-state basis over sixteen years (2002-2018) [6]. This offers a distinctive opportunity for HCI empirical work. On one hand, HCI can conduct natural experiments, comparatively analyzing how technologies and different

laws and their implementations play out differently. On the other hand, because companies and stakeholders have the incentive to comply with the strictest local policies that impact their technologies, HCI can leverage policy elsewhere to promote change in their locale of interest.

Emergent HCI-Policy Trans-Disciplinary Methods. Seeing policy as integral to HCI also allows us to imagine not-yet-existent empirical methods to study system-people-policy interaction.

For example, we see an opportunity for empirical research on successful technology and policy designs. Good designs of technologies, social practices, and policies fade into the background of everyday life, making them more difficult to observe (e.g., via stakeholder interviews). Presently, these success stories can offer valuable lessons, yet are almost entirely absent from HCI literature. A key aim here is to make what is generally invisible to lay people visible, to allow technology designers, policymakers, users, and other stakeholders to reflect on the kinds of system-people-policy interaction design they would find meaningful and actionable.

Next, if HCI is ready to embrace passive, observational empirical work on system-people-policy interaction, are we ready to accept research that creates probes to study tech-policy interactions?

Research through Litigation.

Kirkham [53] is a computer scientist with seven years of experience bringing legal cases to Courts and Tribunals. He does so for the purpose of (1) understanding how the law and legal system operates on the ground and, secondarily, when possible, (2) changing the law. He named this approach "Research through Litigation."

The idea underlying this approach is that the law is a set of ill-defined rules that variably apply to changing real-world circumstances. We cannot understand how a law works simply by reading the law on paper. Further, good laws often fail because of a lack of institutional capacity to implement them.

To understand this gap between law on paper and law on the ground, Kirkham carefully picked legal cases that are good "test cases." The legal process allows him to argue for his case, observe people's reactions, and, in the process, generate voluminous documents (e.g., detailed correspondence within tech companies, expert testimonies) unlikely to be obtained through other means. He then analyzes this process using autoethnography and document analysis methods. This approach has surfaced "*some serious (and somewhat surreal) concerns with the operation of the justice system*" and indeed changed the law on several occasions [53].

People's experience of technology is highly context-dependent. HCI empirical researchers routinely use artifacts (e.g., technology probes, cultural probes) to give users and stakeholders a taste of the future. In so doing, researchers gain better insights into how people might interact with future technologies and derive more informative design implications.

People's experience of policies is also highly context-dependent. The proposal of Research through Litigation highlights this complexity that most prior HCI-policy research neglects. In a sense, the legal cases Kirkham created are analogous to technology probes.

They mean to reveal how law works on the ground, and on occasion, even improve the law. In this light, does this seemingly radical method merit more consideration? Can HCI effectively simulate how courts operate, as a way to assess the effects of new tech and new policy on the ground?

6.2 Integrating Policy into Design Work

HCI-Policy Synergies and Tensions. As mentioned earlier (§2.2), HCI and policy design share many synergies: Both seek to improve technologies' human impact. Both must thoughtfully shape their design processes, because design processes anchor design outcomes. Both must navigate the tensions between design goals and technical feasibility constraints [8, 40]. HCI and policy designers already share formal methods such as participatory design [25, 91].

Seeing policy as integral to HCI's design inquiries illuminates additional connections between them. For example, we realize that both HCI and policy use service design methods [48]. Although service designers do not always mention the word policy, they routinely design computational systems, organizational workflows, and related corporate or platform policies in tandem [33, 82]. When service designers specify a *freemium* model for a new app, they specify policies around who can access what quality of service at what price, a simple micro-economic policy. They shape user experience by innovating system-human-policy interactions.

Other forms of HCI-policy design collaboration are also emerging. For example:

- *Designing public service/policy.* The boundary between designing computational systems for government agencies and shaping public policy is porous. Therefore, many HCI design projects in the public sector involve policy design. For example, researchers redesigned the U.S. Postal Service (USPS) manual and, in doing so, reshaped policies around how small businesses can access postal services, alleviating their legal burdens [49].
- *Designing both computational systems and policies based on a shared strong concept.* For example, FAccT publications on what "fairness" [13] or "representativeness" [21] means catalyzed synergistic designs of AI systems and their regulations.
- *Designing policy evaluation metrics.* For example, Jin [46] collaborated with legal professionals to define evaluation criteria for an algorithm in the U.S. criminal legal system. This design bridged HCI and policy experts' notions of a good algorithm and a good legal system.
- *Designing legally mandated tech design processes.* Some HCI researchers promoted HCI's human-centered design processes to become legally required or recommended technology design processes [24].

Tensions can also arise between technology and policy design, for example, due to documentation differences. HCI and policy communities differ significantly in how they document, disseminate, and attribute their designs. One communicates design ideas via computational artifacts, demos, and pictorials. The other uses policy white papers and law reviews. As a result, social connections (e.g., an HCI designer taking a gap year to work at government agencies) can seem like the only route to simultaneously engage in technology and policy design.

Secondly, coordinating HCI and policy design processes is challenging, because of their multiple and disparate temporal patterns. Innovations in HCI system design can occur rapidly, but may also progress slowly, for example, when developing large computational infrastructures. Policy change can culminate over decades, but can also suddenly accelerate, for example, when public interests surge (e.g., on regulating generative AI [41]) or certain public events occur (e.g., the outbreak of COVID-19 [4]).

Lastly but crucially, power differences complicate synergistic HCI-policy designs. Corporate and public policies can shape computational system designs and user interactions with their authority and power. The reverse is less true.

Low-Hanging Fruits in Addressing HCI-Policy Tensions. By explicating the tensions between HCI and policy design, we can identify practical solutions to address them. For example, improving documentation of HCI-policy design collaborations. Prominent HCI research venues can contribute to this by accepting coordinative tech and policy designs into their proceedings.

Other near-future solutions can further support and popularize the emergent forms of HCI-policy design collaboration listed above. For example, we see an opportunity to extend the practice of designing legally mandated tech design processes into the context of privacy by design. While policy documents broadly call for "privacy by design", it is not always clear *how* that should be done. Consequently, legal assessment and risk management practices ended up operationalizing these design tasks [94]. HCI researchers have already created many design methods to address privacy, and can more explicitly connect their work with the goals of privacy by design as articulated in policy documents. More ambitiously, HCI researchers can actively promote HCI design methods (e.g., user-centered design, value-sensitive design) as equally essential approaches, on par with legal and risk management strategies in privacy by design.

HCI designers can strategically align their tech design goals with policymaking timelines. Upon technical and/or social change, HCI's novel system design contributions can help fill in the temporary policy vacuums [65], bridging the gap between what the law protects and promotes and what people believe the law ought to protect or promote [11]. In doing so, HCI's system designs can help protect people and circumstances that the law fails to protect [8].

Emergent HCI-Policy Trans-Disciplinary Methods. Seeing policy as integral to HCI sets the stage for HCI and policy designers to collaboratively design technologies, social practices, and policies in tandem, creating an artful tech-policy-human interplay. Service designers already do so to an extent, though typically within the confines of one company's technologies and policies [20, 29]. There is an open opportunity to expand the emergent practice of simultaneous tech-and-policy design to larger-scale design contexts.

Sandhaus et al. [75] proposed iCAPS (Integrative Prototyping of City Environments, Autonomous Vehicle Behaviors, and Policies), a simulation-based prototyping platform. Build upon the digital twin of a city, this platform simulates and visualizes how design changes to autonomous vehicles' (AVs') driving behaviors, city design, and related policies might influence city design goals (e.g., road safety,

neighborhood equity, pollution). Via simultaneous AV-city-policy design, this platform promises to catalyze safer, more equitable, and more sustainable future cities [35].

This platform is also a boundary object. AV, city, and policy design decisions belong to many different private corporations and government agencies. Although proven beneficial [72], coordinating these design decisions remains challenging in practice. The iCAPS platform addresses this challenge by bringing AV, city, and policy designers together and moderating their design actions.

This emergent method embodies the idea of designing technology and policy simultaneously, and proposes one concrete way of operationalizing it. It raises at least two sets of useful questions.

This method highlights the power struggles among designers when they design technology and its many related policies in tandem. Systemic thinking is integral to HCI design expertise. It allows designers to grasp how various laws, policies, technologies, and contextual factors can collectively influence people's experiences (Figure 1). However, it would be naive to think that tech designers can re-design related laws and policies without multiple policy experts at the table. Must designers convene *all* relevant policymakers to the same table, in order to design tech and the many policies it involves in tandem? If that is not feasible, which laws and policies should HCI designers prioritize?

This method also underscores the challenges of evaluating system-people-policy interaction design. While HCI design typically iterates and learns from failures, a failed public policy can have irreversible consequences [73]. The proposed prototyping platform addresses this challenge through simulation, computationally predicting societal outcomes for each AV-city-policy design. This contrasts with HCI's traditional approach of evaluating designs with real stakeholders. Do HCI communities support this shift?

6.3 Integrating Policy into HCI Futuring Work

HCI-Policy Synergies and Tensions. Both policy and HCI communities want to foresee how emergent technological capabilities might interact with people and societies, cause disruptions, and create needs for new computational system designs, social practices, and/or policies. To do so, the field of HCI has methods such as speculative design and consequence scanning [3, 50]. Policy actors have methods such as forecasting [76] and the Precautionary Principle [55].

Yet at least two differences between HCI's and policy's approaches can hinder their collaboration. Prior research often criticized HCI's policy engagements being too reactive [43]. These hindrances might explain why.

First is their different views of risks. HCI's speculative designs indicate qualitative risks, often without indicating the level of urgency. Their goal is to provoke discussion and providing cautionary tales for new technologies' adoption and (mis)use. Risk and urgency levels, however, are essential for policy actors' futuring work: Is a new technology's societal harm catastrophic but unlikely (like nuclear war), slowly unfolding but inevitable (like climate change), or so catastrophic and inevitable that policy should forbid it outside of research labs if not entirely (like genetic editing of deadly diseases)? The answer to this question is crucial for policy-making [45], yet fall outside of existing HCI futuring methods.

Second is their different views of design. As we have hinted at earlier in the paper, while speculative design is a process of open-ended knowledge inquiry [94], some legal scholarship sees design as merely the implementation of a given policy requirement. Many policymakers engaged designers only after they had identified a policy goal (e.g., ensure citizens' right to privacy) and had translated it into system requirements (e.g., providing individual control over personal information by increasing consumer notice and choice [34]). They engaged HCI designers only to implement

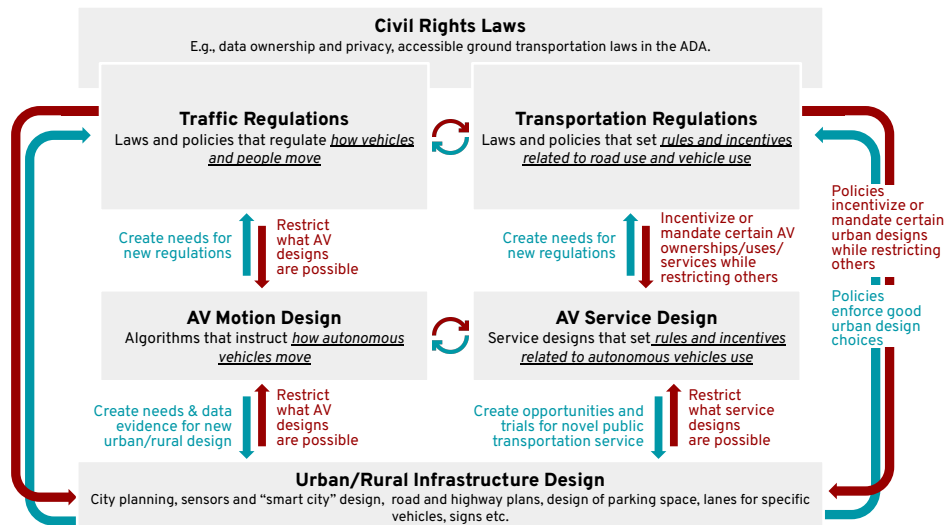


Figure 1: The many ways autonomous vehicle (AV) behavior design, city design, and law and policy interact in the U.S. legal context. They enable (green arrows) and constrain (red arrows) each other [75].

these requirements (e.g., design the standardized cookie opt-out menus [84]).

Low-Hanging Fruits in Addressing HCI-Policy Tensions. By explicating the connections and tensions between HCI and policy’s futuring methods, we start to identify easy ways to strengthen the connections and alleviate tensions.

For example, HCI’s speculative design work can articulate its <Implications for Policy> more explicitly and concretely. What emergent technologies (e.g., facial recognition AI) might require new methods to meet old policy goals? In what ways may current laws and policies fall short in addressing a problem? What values that people hold do not fit into existing legal frameworks (e.g., changing privacy norms)? How might policymakers define a new policy problem space (e.g., “dark patterns” [37]), based on the emergent socio-technical issues this work has identified? Are new policy initiatives necessary [62]?

Explicating policy implications of HCI’s speculative design work can potentially increase its policy impact. In addition, this approach can showcase the value of speculative design throughout the policy-making process, gradually correcting misconceptions about HCI’s design expertise among some policy actors. Imagine, if HCI researchers had defined the new social issues and policy needs that chatGPT would entail, within the four years after its neural architecture was first published in academia and before its public release. These insights might just tackle the challenge of “*bringing policymakers to the table*” [77].

Emergent HCI-Policy Trans-Disciplinary Methods. Seeing policy as integral to HCI’s futuring work also gives us license to imagine entirely new types of futuring research. For example, speculative design practice might embrace speculations around law and policy. It can speculate about the social and technical implications of an emergent policy proposal. It can also envision new uses and forms for emerging technologies, new categories of law and policies, and new social practices. After all, these factors and the interactions among them simultaneously shape emergent technologies’ societal impact.

More ambitiously, we see opportunities in exploring new technological risk forecasting methods that combine HCI’s qualitative, provocation-oriented approaches with policy’s semi-quantitative, action-oriented approaches. To help provoke our imaginations around these opportunities, we describe one such method emergent from our workshop; a method that some may consider too progressive, or too tilted towards policy methods and not enough towards HCI. Presenting this method as a provocation, we ask: Suppose we (HCI communities) are ready to embrace policy as an integral part of its futuring work. Are we also ready to embrace political science’s view of rigor in forecasting?

AI Capability Forecasting:

Dardaman and Gupta [23] proposed to quantitatively forecast AI capability growth, in order to change the reactionary stance HCI researchers, policymakers, and organizations often found themselves in when regulating AI. Such forecasts, and the process of generating and debating them, can help them build mid- and long-term strategies for addressing AI’s societal impact. The quantitative forecasting method has three steps:

1. Define a technology capability prediction problem with clear evaluation criteria;
2. Invite many researchers and forecasters with a variety of relevant expertise to submit their quantitative predictions, along with documentations of their methods, assumptions, and uncertainty measures;
3. Aggregate the forecasts by examining their degree of consensus. As a secondary step, some researchers bring forecasters together, enabling them to challenge each others’ assumptions, sharpen their methods, and collaboratively converge on the most likely predictions.

This method challenges HCI’s assumption that speculations of technology harms are only qualitative. To inform policy actions, analysis of how likely or how soon the harm will become a reality is critical. In order to impact policy, is HCI’s speculative design work open to alternative criteria of rigor in its speculations?

Throughout Chapter 6, we illustrated that seeing policy as integral to HCI’s intellectual endeavors can enhance these endeavors. Examining existing HCI literature (including ones previously considered at the fringes of HCI) through this new lens, we can begin to see inherent tensions between HCI and policy methods, identify pragmatic approaches to alleviate these tensions, and deliberate on emergent methods that address these tensions in thought-provoking ways. Going forward, we encourage HCI communities to together foster a mosaic of distinctively HCI methods and knowledge contributions that blend system, human, and policy expertise to various degrees and in diverse ways, creating robust, grass-root connections between the two fields. Chapter 6 outlines an initial draft of this new landscape.

7 OPPORTUNITIES IN ENHANCING HCI’S COLLECTIVE IMPACT ON POLICY

Now, we turn our attention to HCI-policy collaboration at a community level. We previously proposed that, by seeing policy as integral to their intellectual concerns, HCI communities will coordinate their diverse policy engagements, thereby amplifying HCI’s collective impact on policy outcomes. In this chapter, we outline four emergent opportunities for such coordination by synthesising existing HCI and policy literature.

7.1 When (Not) to Policy

We envision a future where HCI communities strategically choose when and when not to use policy to address computation’s societal issues. A crucial step toward realizing this vision is establishing a principled understanding of policy’s affordances and limitations, compared with other HCI’s tools. There is an under-explored area in HCI research and a clear opportunity for future work.

On the one hand, HCI can further exploit the affordances of policy. For example, current HCI-policy work focused overwhelmingly on regulating technology harms. Less discussed is how policy can also enable new actions. For example, copyright law attempts to balance protecting creators from intellectual property harms, while also encouraging new innovations for social good through fair use. How might HCI leverage policy to promote the positive impact of technologies?

On the other hand, policy interventions are not a cure-all, yet their limitations are overlooked in current HCI literature. They can be prone to their own forms of inequities and bias [9]. Moreover, law and policy solutions promote incremental changes within an existing legal and institutional structure. They typically do not “change the system,” particularly not overnight [17]. Depending on the degree and kind of social change HCI hopes to create, sometimes, other approaches (e.g., community organizing) might be more effective than policy interventions.

Explicating the limitations of policy can also reveal new opportunities for technology design. For example, HCI researchers often criticize that capitalist goals drive many technology product designs and cause ethical harm. They then seek remedies from policy interventions. However, law and policy can be just as enwrapped in promoting or serving institutions of capitalism [22]. In order to improve technologies’ societal impact, HCI sometimes needs to consider designing technologies that help the disenfranchised populations fight unjust policies [10].

7.2 Which Mode of Policy Engagement & When

We envision a future where HCI communities will wisely choose the extents and modes of policy engagement, based on the specific technology and social issue at hand. A crucial step toward realizing this vision is to develop a principled understanding of *how* the most productive approach to policy engagement differs according to context. For example, when is appending a thoughtful <Implications for Policy> to a classic HCI empirical study already sufficient? In what contexts is establishing deep, ongoing partnerships with policymakers not just beneficial, but necessary?

The history of HCI—more specifically, the three waves of HCI [15]—can serve as an initial scaffolding for this discussion. This history seems to suggest that as a computing system becomes more deeply and broadly integrated into societal processes, addressing its impact requires deeper collaboration between HCI and policy.

- The first wave of HCI focuses on making computers more efficient and functional for computer experts. These efforts involved minimal policy considerations.
- The second wave of HCI focuses on making personal computers easier to use and more accessible to everyday users. These efforts catalyzed the ISO usability standards and accessibility laws, which in turn enhanced later HCI work. Here, HCI and policy work operated asynchronously. Appending a thoughtful <Implications for Policy> to an HCI empirical study on usability or accessibility could go a long way.
- The third wave of HCI focused on weaving computing into daily life through smartphones, tablets, and wearable devices. Around this time, HCI privacy scholars like Lorrie Cranor drafted national privacy policies at the FTC and designed standardized cookie opt-out menus [24]. State laws later recommended these designs [84]. Here, while HCI and policy work remained somewhat asynchronous, technology and policy design processes started to overlap.

Do the three waves of HCI effectively stratify technologies’ needs for HCI-policy collaborative work? Is there a need for additional or more fine-grained categories? These questions provide fertile ground for future research.

For example, the societal impact of interactive cyber-infrastructure appears to demand particularly close HCI-policy partnerships [30], and therefore might merit its own category. Interactive cyber-infrastructure refers to interactive systems that not only host, *but govern* various interactions on a societal-level scale². For instance, social media platforms whose algorithms govern online discourse, gig work platforms whose algorithms govern work assignments, pervasive sensing networks that govern how smart cities operate, AI Foundation Models (FMs) such as chatGPT that govern downstream AI innovations, and virtual reality (VR) “*universes*” where governments offer public services [7]. Understanding and improving the societal impact of cyber-infrastructure necessitate close HCI-policy partnerships, for three reasons.

- Because of the immense scale of these systems, HCI’s conventional empirical methods may struggle to trace or prove their societal impact. Law and policy methods (e.g., Research through Litigation) might help;
- Because of the immense power the cyber-infrastructure owners wield (think the power of OpenAI), HCI designers heavily rely on policies that mandate system transparency and/or access, in order to do human-centered design work. Designing one more socially beneficial GPT prompt is simply not as impactful as devising a way of reducing the creation of harmful prompts. The latter necessitates policy interventions.
- Because these systems are infrastructures, they shape how HCI researchers and practitioners work. *Should CHI keep publishing novel GPT applications, even though they might be soon regulated out of existence?* Questions like this highlight the necessity of HCI-policy considerations from the outset.

7.3 How to Maximize HCI’s Collective Impact

We envision a future where HCI can assemble the different subsets of its work to strategically influence different stages of policy-making, and to inform different policy actors [48, 52]. When sudden momentum for policy change appears, HCI communities can quickly assemble their diverse relevant work to seize the opportunity.

Preparing the ground for policy change. Even policy changes that seem to occur “suddenly” have longer histories of community discussion and debate. This period is an opportune time for cultivating relationships with policy actors, engaging in policy discussions that occur in the background of everyday life, and incubating HCI empirical and design work.

HCI researchers and practitioners can become more actively engaged in everyday policy discussions, for example, by responding to regulatory agencies’ requests for public comment [1] and writing Op-Ed articles (and other forms of tech journalism). In these discussions, HCI professionals can share metaphors that help policymakers understand the affordance of emergent technologies, offer empirical evidence of technological harm as expert evidence for future policies, recommend human-centered design methods to be legally mandated or recommended, and propose value-sensitive

²To be clear, here we use the term “*interactive cyber-infrastructure*” to denote a distinct category of computational systems, rather than as an analytical lens on the sociotechnical systems that shape society. We use the word “*infrastructure*” differently than in Science and Technology Studies (STS) and cultural studies (e.g., [30, 31, 79]).

metrics for evaluating future policies. All this work can prepare the ground for future policy change.

In addition, the potentially long “incubation” period of policy change is also opportune for long-term HCI research, such as conducting participatory design workshops and comparatively analyzing how different tech-and-policy designs play out in different locales.

Catalyzing a groundswell of public interest and policy demand. Law and policy can change quickly in response to current events (e.g., COVID-19 [4]), technical developments (e.g., chat-GPT [90]), and other changes. Yet that does not mean HCI has to wait passively for a pandemic of opportunity for policy change open up.

HCI’s speculative, provocative, and participatory design work can catalyze a groundswell of public opinion around HCI issues, creating a new policy demand. Critical HCI is already somewhat successful in this regard [91]. Additional opportunities exist for HCI to engage interests groups, corporations, and policymakers more pointedly. The aim is to align the stars: When “*a problem is recognized, a solution is available, and the political climate happens to be right*”, a window of opportunity for substantial policy change opens [52].

Seizing sudden momentum of policy change. When a window of opportunity for policy change opens up, the heightened levels of attention are fleeting [52]. HCI communities should be prepared to seize the political momentum.

HCI researchers can prepare rapid policy responses by re-using their existing work, but adding a “hook” that connects it to the current events and political momentum. <Implications for Policy> sections of past HCI futuring work, literature reviews summarizing existing empirical evidence, and policy proposals tested in prior participatory workshops are all valuable, reusable resources here. These responses can take the form of op-eds, memos, policy white papers, public agency comments, and more.

7.4 How Might HCI Institutions Help

Finally and importantly, we call for HCI institutions to nurture the ties between HCI and policy communities and to amplify HCI’s collective voice in the policy world.

Including policy education in HCI; including HCI in policy education. We envision a future where policy-making and policy implementation become a standard element of HCI education. This education can cultivate future HCI researchers and practitioners who can navigate layers of government agencies, navigate laws of intricately overlapping jurisdictions, connect with policymakers, and integrate policy expertise into their own work. We also see an opportunity for HCI to become part of legal and political science education. Today, many leading law schools already offer curriculum on AI ethics, as do many HCI programs. Harnessing and expanding these existing connections incubate future HCI-policy collaborations of all sorts.

Celebrating small wins. From accumulating research evidence, to building relationships with policymakers, to acquiring policy expertise, HCI researchers and practitioners’ journey to policy impact is arduous. Moreover, such impacts are often invisible: Citing academic references or attributing individual researchers is not

a convention in legal or public policy documents. How can HCI institutions (such as CHI and CSCW) acknowledge and celebrate small wins in HCI researchers’ policy engagement?

Consider: Can CHI proceedings accept insightful literature reviews that synthesize prior HCI works’ <Implications for Policy> into actionable policy proposals? Can CHI offer a short paper track where researchers share their responses to regulatory agencies’ calls for comments? When public agencies or policymakers call for scientific evidence, they often require a response within days or weeks [24]. In parallel to peer-reviewed publications, can HCI conferences create channels for faster-paced policy discourses and contributions? Such small recognitions and incentives have the potential to significantly boost policy engagements, particularly for early-career professionals in HCI.

Amplifying HCI’s collective voice in the policy realm. HCI institutions can help strengthen the collective ties between HCI and policy communities. For example, the registration costs for HCI conferences may be difficult for public servants to justify; are there ways to make the HCI community more accessible to these policy audiences? Can CHI sets up booths for policymakers and public servants (just like it does for recruiters), helping HCI researchers connect with them?

HCI institutions can help amplify HCI communities’ collective voices in the policy realm. For example, HCI can offer social infrastructures that enable HCI professionals to collaborate on public comments, statements, and policy recommendations. This approach can help more evenly distribute the policy engagement work across community members, and can amplify the collective voice of HCI communities.

8 CLOSING REMARKS

The field of HCI has a history of self-improvement: identifying a problem, creating a new way of working, and then using evidence of success to argue for a change in the community. Early computer systems caused chaos when they first entered workplaces, facing users who were not computer scientists. Early HCI researchers (before even calling themselves HCI) created a new way of designing technologies in response, that is, *user-centered design*. When computing moved out of the workplace and into people’s everyday lives, *users* became *consumers*, and they had choices. In response, HCI practitioners worked to improve the situated experience of technologies, and researchers dug in and unpacked the very concept of “*experience*”.

All this bodes well for the changes facing our community today. In current technology discourses, people are not only users and consumers. They are members of communities, citizens, and participants in lived democracy. Technologies are not only systems. They are socio-technical platforms and cyber-infrastructure that govern societal processes. People’s “*experience*” of technologies goes beyond efficiency or pleasure, but concerns the future of work, equity, sustainability, and more. Within such a milieu, we argue for a change in our community. We argue that HCI should re-examine the relationship among systems, people, and policy, and reflect on how HCI wants to position itself in this relationship. This paper provides an additional perspective to this unfolding dialogue.

ACKNOWLEDGMENTS

Qian Yang's effort to help organize and participate in the workshop was partially supported by National Science Foundation under grants IIS-2212431, IIS-2313077, and Schmidt Futures' AI2050 Early Career Fellowship. Richmond Y. Wong's effort was supported by a Georgia Tech Ivan Allen College Small Grant for Research (SGR). Sabine Junginger's work on integrating policy and service design is supported by the Swiss National Science Foundation SINERGIA CRSII5 189955 VA-PEPR (Voice Assistants, People Experiences, Practices and Routines). 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.

REFERENCES

- [1] 2011. https://www.federalregister.gov/uploads/2011/01/the_rulemaking_process.pdf
- [2] 2012. United States v. Jones, 565 U.S. 400 (2012). <https://supreme.justia.com/cases/federal/us/565/400/>
- [3] 2017. <https://www.gov.uk/guidance/open-policy-making-toolkit>
- [4] 2020. <https://csrreports.congress.gov/product/pdf/LSB/LSB10511>
- [5] 2021. https://www.ftc.gov/system/files/documents/public_events/1586943/dark_patterns-workshop-bios.pdf
- [6] 2022. <https://www.ncsl.org/technology-and-communication/security-breach-notification-laws>
- [7] Accenture. 2023. Government enters the metaverse. <https://www.accenture.com/us-en/insightsnew/us-federal-government/technology-vision-2022>
- [8] Mark S Ackerman. 2000. The intellectual challenge of CSCW: the gap between social requirements and technical feasibility. *Human-Computer Interaction* 15, 2-3 (2000), 179–203.
- [9] Kendra Albert and James Grimmelmann. 2023. Do the Right Thing. *Commun. ACM* 66, 5 (2023), 18–20.
- [10] Tran Anh. 2023. Adversarial Engagements Between Technology and Policy. In *Workshop on Designing Technology and Policy Simultaneously: Towards A Research Agenda and New Practice*.
- [11] Kenneth A Bamberger and Deirdre K Mulligan. 2011. Privacy on the Books and on the Ground. *Stanford Law Review* (2011), 247–315.
- [12] 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>
- [13] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2017. Fairness in machine learning. *Nips tutorial* 1 (2017), 2017.
- [14] Eli Blevins. 2007. Sustainable interaction design: invention & disposal, renewal & reuse. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 503–512.
- [15] Susanne Bødker. 2015. Third-wave HCI, 10 years later—participation and sharing. *interactions* 22, 5 (2015), 24–31.
- [16] 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 Doumouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kavin 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 Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kudithipudi, 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, Avnika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, 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]
- [17] Karrigan S Bork. 2019. An evolutionary theory of administrative law. *SMU L. Rev.* 72 (2019), 81.
- [18] Jed R. Brubaker and Vanessa Callison-Burch. 2016. Legacy Contact: Designing and Implementing Post-Mortem Stewardship at Facebook. In *Proceedings of the 19th CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 2908–2919. <https://doi.org/10.1145/2858036.2858254>
- [19] Matt Burgess. 2023. Criminals Have Created Their Own ChatGPT Clones. <https://www.wired.com/story/chatgpt-scams-fraudgpt-wormgpt-crime/>
- [20] 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>
- [21] Kyla Chasalow and Karen Levy. 2021. Representativeness in statistics, politics, and machine learning. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 77–89.
- [22] Ernesto Dal Bó. 2006. Regulatory capture: A review. *Oxford review of economic policy* 22, 2 (2006), 203–225.
- [23] Emily Dardaman and Abhishek Gupta. 2023. Asking Better Questions – The Art and Science of Forecasting: A mechanism for truer answers to high-stakes questions. In *Workshop on Designing Technology and Policy Simultaneously: Towards A Research Agenda and New Practice*. arXiv:2303.18006 [cs.CY]
- [24] Janet Davis, Harry Hochheiser, Juan Pablo Hourcade, Jeff Johnson, Lisa Nathan, and Janice Tsai. 2012. Occupy CHI Engaging U.S. Policymakers. In *CHI '12 Extended Abstracts on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI EA '12). Association for Computing Machinery, New York, NY, USA, 1139–1142. <https://doi.org/10.1145/2212776.2212406> Note: Panel slides can be found at <https://web.archive.org/web/20170829055308/http://www.cs.grinnell.edu/~davisjan/chi-us-public-policy/chi2012-panel.pdf>. Accessed: 2023-07-31.
- [25] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–23.
- [26] Carl DiSalvo, Jonathan Lukens, Thomas Lodato, Tom Jenkins, and Tanyoung Kim. 2014. Making public things: how HCI design can express matters of concern. In *Proceedings of the SIGCHI Conference on Human factors in Computing Systems*. 2397–2406.
- [27] Lynn Dombrowski, Ellie Harmon, and Sarah Fox. 2016. Social Justice-Oriented Interaction Design: Outlining Key Design Strategies and Commitments. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems* (Brisbane, QLD, Australia) (DIS '16). Association for Computing Machinery, New York, NY, USA, 656–671. <https://doi.org/10.1145/2901790.2901861>
- [28] Kees Dorst and Nigel Cross. 2001. Creativity in the design process: co-evolution of problem–solution. *Design studies* 22, 5 (2001), 425–437.
- [29] 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>
- [30] Paul N Edwards, Steven J Jackson, Geoffrey C Bowker, and Cory Philip Knobel. 2007. Understanding infrastructure: Dynamics, tensions, and design. (2007).
- [31] W. Keith Edwards, Mark W. Newman, and Erika Shehan Poole. 2010. The Infrastructure Problem in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '10). Association for Computing Machinery, New York, NY, USA, 423–432. <https://doi.org/10.1145/1753326.1753390>
- [32] Jodi Forlizzi. 2018. Moving beyond user-centered design. *interactions* 25, 5 (2018), 22–23.
- [33] Jodi Forlizzi. 2018. Moving beyond user-centered design. *interactions* 25, 5 (2018), 22–23.
- [34] Robert Gellman. 2022. Fair Information Practices: A Basic History-Version 2.22. Available at SSRN (2022).
- [35] 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*.
- [36] Colin M Gray. 2022. Expert Report of Colin M. Gray, Ph. D. on Dark Patterns (Public Redacted Version). *STATE OF ARIZONA, ex rel. MARK BRNOVICH, Attorney General, Plaintiff, v. Google LLC* (2022). <https://www.azag.gov/sites/default/files/2022-09/Expert%20Report%20of%20Colin%20M.%20Gray%2C%20Ph.D..pdf>
- [37] Colin M Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L Toombs. 2018. The dark (patterns) side of UX design. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–14.
- [38] Louise Møller Haase and Linda Nhu Laursen. 2019. Meaning frames: The structure of problem frames and solution frames. *Design Issues* 35, 3 (2019), 20–34.

- [39] Oliver L Haimson, Daniel Delmonaco, Peipei Nie, and Andrea Wegner. 2021. Disproportionate removals and differing content moderation experiences for conservative, transgender, and black social media users: Marginalization and moderation gray areas. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–35.
- [40] Michael T Hayes. 2002. *The limits of policy change: Incrementalism, worldview, and the rule of law*. Georgetown University Press.
- [41] David Hickton. 2023. Chatgpt could transform society and its risks require quick regulation. <https://thehill.com/opinion/technology/3812597-chatgpt-could-transform-society-and-its-risks-require-quick-regulation/>
- [42] Jason Hong. 2023. What Can the FATE Community Learn from the Successes and Failures in Privacy? <https://cacm.acm.org/blogs/blog-cacm/271991-what-can-the-fate-community-learn-from-the-successes-and-failures-in-privacy/fulltext>. Accessed: 2023-07-29.
- [43] 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>
- [44] Erich Jantsch. 1972. Inter- and transdisciplinary university: A systems approach to education and innovation. *Higher education* 1, 1 (1972), 7–37.
- [45] Sheila Jasanoff. 2014. Serviceable truths: Science for action in law and policy. *Tex. L. Rev.* 93 (2014), 1723.
- [46] Angela Jin. 2023. Opportunities for Coordinating Policy and Technology Design for Algorithmic Accountability: Insights from the U.S. Criminal Legal System. In *Workshop on Designing Technology and Policy Simultaneously: Towards A Research Agenda and New Practice*.
- [47] 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>
- [48] Sabine Junginger. 2016. *Transforming Public Services by Design: Re-orienting policies, organizations and services around people*. Taylor & Francis.
- [49] Sabine Junginger. 2016. The USPS domestic mail manual transformation project. In *Transforming Public Services by Design: Re-orienting policies, organizations and services around people*. Taylor & Francis.
- [50] Rob Katz. 2021. How to run a consequence Scanning workshop. <https://medium.com/salesforce-ux/how-to-run-a-consequence-scanning-workshop-4b14792ea987>
- [51] Sarah Kessler. 2013. Uber: When innovation outpaces the law. <https://www.fastcompany.com/3001169/uber-when-innovation-outpaces-law>
- [52] John W Kingdon and Eric Stano. 1984. *Agendas, alternatives, and public policies*. Vol. 45. Little, Brown Boston.
- [53] Reuben Kirkham. 2023. (Legal Design) Research through Litigation. In *Workshop on Designing Technology and Policy Simultaneously: Towards A Research Agenda and New Practice*. arXiv:2303.14336 [cs.HC]
- [54] Julie Thompson Klein. 1990. *Interdisciplinarity: History, theory, and practice*. Wayne state university press.
- [55] David Kriebel, Joel Tickner, Paul Epstein, John Lemons, Richard Levins, Edward L Loecher, Margaret Quinn, Ruthann Rudel, Ted Schettler, and Michael Stoto. 2001. The precautionary principle in environmental science. *Environmental health perspectives* 109, 9 (2001), 871–876.
- [56] Benjamin Laufer, Sameer Jain, A Feder Cooper, Jon Kleinberg, and Hoda Heidari. 2022. Four years of FAccT: A reflexive, mixed-methods analysis of research contributions, shortcomings, and future prospects. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 401–426.
- [57] Jonathan Lazar. 2010. Interacting with Public Policy. *Interactions* 17, 1 (jan 2010), 40–43. <https://doi.org/10.1145/1649475.1649485>
- [58] 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>
- [59] Jonathan Lazar. 2017. Let's strengthen the HCI community by taking a gap year! *Interactions* 25, 1 (2017), 20–21.
- [60] Jonathan Lazar, Julio Abascal, Janet Davis, Vanessa Evers, Jan Gulliksen, Joaquim Jorge, Tom McEwan, Fabio Paternò, Hans Persson, Raquel Prates, Hans von Axelson, Marco Winckler, and Volker Wulf. 2012. HCI Public Policy Activities in 2012: A 10-Country Discussion. *Interactions* 19, 3 (may 2012), 78–81. <https://doi.org/10.1145/2168931.2168947>
- [61] Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. 2015. Working with machines: The impact of algorithmic and data-driven management on human workers. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 1603–1612.
- [62] Joseph Galen Lindley, Paul Coulton, Haider Akmal, and Brandin Hanson Knowles. 2017. Anticipating GDPR in Smart Homes Through Fictional Conversational Objects. (2017).
- [63] Michael A. Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. 2020. Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376445>
- [64] Lisa Messeri and Janet Vertesi. 2015. The Greatest Missions Never Flown: Anticipatory Discourse and the “Projectory” in Technological Communities. *Technology and Culture* (2015), 54–85.
- [65] James H Moor. 1985. What is computer ethics? *Metaphilosophy* 16, 4 (1985), 266–275.
- [66] Lisa P. Nathan and Batya Friedman. 2010. Interacting with Policy in a Political World: Reflections from the Voices from the Rwanda Tribunal Project. *Interactions* 17, 5 (sep 2010), 56–59. <https://doi.org/10.1145/1836216.1836231>
- [67] Andy Newman. 2019. I found work on an Amazon website. I made 97 cents an hour. *The New York Times* 15 (2019).
- [68] Hope Nudzur. 2009. What is “policy”, a problem: solving definition or a process conceptualisation. *Educational futures* 2, 1 (2009), 85–96.
- [69] Billy Perrigo. 2023. Exclusive: Openai lobbied E.U. to water down AI regulation. <https://time.com/6288245/openai-eu-lobbying-ai-act/>
- [70] B Guy Peters. 2021. *Advanced introduction to public policy*. Edward Elgar Publishing.
- [71] Michael Ramesh. 2003. *Studying public policy: Policy cycles and policy subsystems*. Don Mills, Ont.: Oxford University Press.
- [72] William Riggs, Melissa Ruhl, Caroline Rodier, and Will Baumgardner. 2019. Designing Streets for Autonomous Vehicles. In *Road Vehicle Automation 6*. Springer International Publishing, 111–122. https://doi.org/10.1007/978-3-030-22933-7_12
- [73] Horst W. J. Rittel and Melvin M. Webber. 1973. Dilemmas in a General theory of planning. *Policy Sciences* 4, 2 (June 1973), 155–169. <https://doi.org/10.1007/BF01405730>
- [74] Virpi Roto, Jung-Joo Lee, Effie Lai-Chong Law, and John Zimmerman. 2021. The overlaps and boundaries between service design and user experience design. In *Designing Interactive Systems Conference 2021*. 1915–1926.
- [75] Hauke Sandhaus, Wendy Ju, and Qian Yang. 2023. Towards Prototyping Driverless Vehicle Behaviors, City Design, and Policies Simultaneously. In *Workshop on Designing Technology and Policy Simultaneously: Towards A Research Agenda and New Practice*. arXiv:2304.06639 [cs.HC]
- [76] Tamara Savage. 2023. Improving Technology Forecasting by Including Policy, Economic, and Social Factors. (9 2023). <https://doi.org/10.1145/124103998.v1>
- [77] Anne Spaa, Abigail Durrant, Chris Elsdén, 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>
- [78] Anne Spaa, 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>
- [79] Susan Leigh Star and Karen Ruhleder. 1994. Steps towards an ecology of infrastructure: complex problems in design and access for large-scale collaborative systems. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work*. 253–264.
- [80] Stephanie B Steinhardt and Steven J Jackson. 2014. Reconciling rhythms: plans and temporal alignment in collaborative scientific work. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. 134–145.
- [81] Stephanie B Steinhardt and Steven J Jackson. 2015. Anticipation work: Cultivating vision in collective practice. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*. 443–453.
- [82] Marc Stickdorn and Jakob Schneider. 2012. *This is service design thinking: Basics, tools, cases*. John Wiley & Sons.
- [83] Montreal The Global Partnership of Artificial Intelligence Innovation Workshop. 2023. *Policy Brief: Generative AI, Jobs, and Policy Response*. Technical Report. The Global Partnership of Artificial Intelligence. <https://gpai.ai/projects/future-of-work/policy-brief-generative-ai-jobs-and-policy-response-innovation-workshop-montreal-2023.pdf>
- [84] Daniel Tkacik. 2020. CyLab researchers design privacy icon to be used by California law. <https://www.cylab.cmu.edu/news/2020/12/11-donotsell.html>
- [85] Carlos Toxtli, Siddharth Suri, and Saiph Savage. 2021. Quantifying the invisible labor in crowd work. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–26.
- [86] Lachlan Urquhart and Tom Rodden. 2017. New directions in information technology law: learning from human-computer interaction. *International Review of Law, Computers & Technology* 31, 2 (2017), 150–169.
- [87] Lachlan D. Urquhart, Glenn McGarry, and Andy Crabtree. 2022. Legal Provocations for HCI in the Design and Development of Trustworthy Autonomous Systems. In *Nordic Human-Computer Interaction Conference* (Aarhus, Denmark) (NordicCHI '22). Association for Computing Machinery, New York, NY, USA, Article 75, 12 pages. <https://doi.org/10.1145/3546155.3546690>
- [88] Synthetic Users. 2023. Synthetic Users: User Research without the Headaches. <https://www.syntheticusers.com/>
- [89] Merlijn Van Hulst and Dvora Yanow. 2016. From policy “frames” to “framing” theorizing a more dynamic, political approach. *The American review of public administration* 46, 1 (2016), 92–112.

- [90] Tom Wheeler, Niam Yaraghi Nicol Turner Lee, John Villasenor, and Nicol Turner Lee Norman Eisen. 2023. Key enforcement issues of the AI act should lead EU trilogue debate. <https://www.brookings.edu/articles/key-enforcement-issues-of-the-ai-act-should-lead-eu-trilogue-debate/>
- [91] Cedric Deslandes Whitney, Teresa Naval, Elizabeth Quepons, Simrandeep Singh, Steven R Rick, and Lilly Irani. 2021. HCI Tactics for Politics from Below: Meeting the Challenges of Smart Cities. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 297, 15 pages. <https://doi.org/10.1145/3411764.3445314>
- [92] Kyle Wiggers. 2023. The current legal cases against Generative AI are just the beginning. <https://techcrunch.com/2023/01/27/the-current-legal-cases-against-generative-ai-are-just-the-beginning/>
- [93] Richmond Y Wong and Steven J Jackson. 2015. Wireless visions: Infrastructure, imagination, and US spectrum policy. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 105–115.
- [94] Richmond Y Wong and Deirdre K Mulligan. 2019. Bringing design to the privacy table: Broadening “design” in “privacy by design” through the lens of HCI. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, 1–17.
- [95] 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>
- [96] Qian Yang, Richmond Y. Wong, Thomas Gilbert, Margaret D. Hagan, Steven Jackson, Sabine Junginger, and John Zimmerman. 2023. *Designing Technology and Policy Simultaneously: Towards A Research Agenda and New Practice*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3544549.3573827>
- [97] Meg Young, Lassana Magassa, and Batya Friedman. 2019. Toward inclusive tech policy design: a method for underrepresented voices to strengthen tech policy documents. *Ethics and Information Technology* 21 (2019), 89–103.
- [98] JD Zamfirescu-Pereira, Richmond Y Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny can't prompt: how non-AI experts try (and fail) to design LLM prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–21.

APPENDIX

A DATA ANALYSIS PROCESS

We started by proposing many possible problem-solution frames that meet the following three criteria. Next, we iteratively critiqued and improved these emergent frames, also based on these criteria.

- *Effective problem framing*: It can help explain the disparate success levels observed in HCI-policy collaborations so far;
- *Useful solution framing*: It can reveal novel strategies for addressing the disparity;
- *Adaptive solution framing*: The strategies it reveals are flexible, allowing HCI researchers and practitioners to derive specific actions in diverse situations. This flexibility is crucial, because effective approaches to HCI-policy collaboration vary based on contexts [38]. Our goal is not to prescribe fixed methods, but to identify new *avenues* of improvements.

Besides meeting the above criteria, the final frame is also a useful seed for HCI-community-wide discussions. Even the authors lack consensus on whether this problem-solution frame might be too progressive and controversial, or too obvious that it represents what HCI needs to do. In this sense, we see this frame as an invitation (and potentially provocation) for fellow HCI researchers, practitioners, and educators to debate the nature of HCI-policy collaboration challenges. We encourage them to share their frames and collectively deliberate how to best move forward.

B WHY INTERACTIVE, INFRASTRUCTURE TECHNOLOGIES NECESSITATE DEEP, CONTINUOUS HCI-POLICY PARTNERSHIP

AI Foundation Models (FMs) such as chatGPT are telling examples. FMs are machine learning models that AI engineers and end users

alike can easily adapt and modify to create bespoke text-, image-, or video-generation applications [16, 98]. Without policy expertise, HCI's futuring and empirical methods cannot effectively understand FMs' societal impact or collect empirical evidence proving this impact.

- **Each FM powers an AI-model-and-app ecosystem. HCI cannot understand the FM's full impact without considering the policies that govern it.** For example, chatGPT. Its societal impact depends on the various fine-tuned models and downstream apps people build with it [19]. Merely empirically studying how one model (e.g., chatGPT, August 8th, 2023 version) impacts one user here and now misses the forest for the trees. Instead, HCI empirical work must also understand how policies incentivize and regulate the chatGPT-derived-model-and-app ecosystem to understand its societal impact (the “forest.”)

This need for policy considerations is not a coincidence. From sensor networks covering a smart city to algorithm-mediated social media platforms, infrastructure technologies operate at societal-level scales and power an ecosystem of various downstream applications. To understand their societal impact, one must first understand the policies that govern the ecosystem.

- **HCI cannot prove the relationship between FM design choices and societal impact, without HCI and policy join forces.** Companies like OpenAI, which own FMs and their ecosystems, hold enormous power over who can know which aspects of FMs' inner workings and data practices. However, this information is critical for studying FM's legality and social impact. Without the joint forces of policy mandates and HCI knowledge (e.g., on what information is critical for understanding FM's fairness), these companies may never provide the public *meaningful* access to FMs' inner workings.

This need for policy mandates on technology transparency is not coincidental either. The effect of cyberinfrastructure on people is often indirect and invisible. Understanding this effect requires information about its inner workings. However, organizations that own infrastructure technologies (e.g., social media giants, governments) are few and powerful, holding enormous power in guarding this information. In this context, HCI needs policy's power and policy expertise, which can enforce information disclosure and know how to do so without harming tech industry competition.

- **Given that FMs' societal impact is far-reaching and demands quick regulation, empirical researchers need to analyze their societal impact with policymakers side-by-side.**

Unlike in some previous cases, HCI does not seem to have the luxury of time to amass empirical evidence of FMs' societal pros and cons before policymakers demand them [41].

The fact that FMs seem to have impacted “everything everywhere all at once” is not coincidental. From Uber (mobile computing combined with gig work) to social media feed-ranking algorithms, infrastructure technologies, once established and adopted at scale, can directly impact human and human-computer interactions of all sorts. Understanding their societal impact requires HCI and policy to collaborate synchronously. Moreover, they need to

develop such understanding within specific windows of opportunity, for example, before the technology is widely adopted and its associated human practices settle in.

For HCI designers who innovate human-centered systems, policy considerations are also crucial from the outset.

- **If the FM violates the law, HCI's novel FM apps may also violate the law.** By creating one accessible system design, HCI incrementally improves web accessibility. Creating one new GPT app without policy considerations, however, HCI risks violating copyright laws.
- **FMs' unprecedented malleability means that HCI's well-intentioned app designs can easily be misused.** It seems naive to think an FM tool that predicts the likely outcomes of a user study protocol would not have been used to power "*user studies without users*" [88], or an FM tool that helps Reddit moderators to predict upcoming fake news would not have been used to generate fake news instead. Without policy guardrails to regulate bad actors, HCI's well-intentioned app designs can play an unintended role in FMs' societal harms. These seemingly unusual risks of HCI designing illegal or unethical things are not coincidental. Just like VR platforms transformed

how governments offer public services [7] and GPT changed how HCI software designers design apps, infrastructure technologies change how almost everyone—including HCI researchers, practitioners, policymakers, and policy enforcement agencies—works. These changes complicate HCI's efforts to improve or regulate infrastructure technologies' societal impact, offering yet another reason for HCI to collaborate with policy experts.

C RESPONDING TO REGULATORY AGENCIES' REQUESTS FOR COMMENT

In the U.S., regulatory agencies (they are responsible for creating rules that enact laws into practice, such as the FTC or National Institute of Standards and Technology) continually seek input and feedback from both experts and the general public, often through requests for public comments for 30-60 days, allowing the public to comment or submit data related to proposed rules or decision-making [1]. HCI professionals can submit empirical evidence of tech's human impact as expert evidence. This evidence could inform policies that both regulate (or promote) certain forms of technology development and use, or that regulate (or promote certain human behaviors related to technology use.