Putting the ‘Practice’ in Critical Technical Practice: A Decision-Making Pedagogy for the Public Interest Technology Clinic

LAUREN MARIETTA CHAMBERS and DIAG DAVENPORT, University of California, Berkeley, US

FAccT scholarship has sharpened our ability to diagnose harms in sociotechnical systems, yet civic and public interest data work confronts practitioners with a problem that critique alone does not resolve: no analysis is unimpeachable. Progress requires countless decisions: what to represent, how to frame uncertainty, when to stop iterating, and how to remain accountable. How do we prepare students for this reality once work leaves the classroom? We argue that Public Interest Technology (PIT) clinics—semester-long, client-facing, project-based courses working with public interest organizations—provide a distinctive pedagogy for *critical technical practice* because they make such decisions unavoidable and consequential. Drawing on a collaborative instructor autoethnography of a master’s-level PIT clinic at UC Berkeley, complemented by conversations with instructors at peer institutions and analysis of course artifacts, we develop a decision-centered account of clinic learning. We introduce two decision-making mental models that clinics can cultivate: *proximate* project judgment (knowing when to move on while owning the work) and *distal* project judgment (contributing without control). From these, we distill six pedagogical design principles that operationalize student decision-making under constraint, and we analyze the resulting promise, challenges, and risks of clinics with respect to developing adept critical technical practitioners. This work offers a vocabulary and design logic for teaching not only the critique, but the practice of making sociotechnical decisions of practical consequence.

CCS Concepts: • **Social and professional topics → Computing education**; **Model curricula**; Project and people management; *Computing profession*; • **Applied computing** → *Computers in other domains*; • **Human-centered computing**;

Additional Key Words and Phrases: public interest technology clinic, CIS service learning, liberatory pedagogy, decision-making

ACM Reference Format:
Lauren Marietta Chambers and Diag Davenport. 2026. Putting the ‘Practice’ in Critical Technical Practice: A Decision-Making Pedagogy for the Public Interest Technology Clinic. In *The 2026 ACM Conference on Fairness, Accountability, and Transparency (FAccT ’26), June 25–28, 2026, Montreal, QC, Canada.* ACM, New York, NY, USA, 21 pages. https://doi.org/10.1145/3805689.3806542

1 Introduction

In the fall of 2020, the American Civil Liberties Union of Massachusetts partnered with the Spark! Lab[1] at Boston University – an experiential learning lab within BU’s Faculty of Computing and Data Sciences. ACLU-MA was just one of many local organizations serving as public interest clients and providing data projects for an undergraduate course that semester. The civil rights group’s project involved student analysis of data on the Boston Police Department’s overtime practices: policies, expenditures, and hours. Over the course of a few months, an effective and dedicated student team produced a compelling suite of visualizations motivating questions and concerns around police overtime usage. Their efforts paid off in a big way: the students’ findings were packaged by ACLU-MA and presented as written testimony to the Boston City Council via a comprehensive technical report [4]. As

[1]See https://www.bu.edu/spark/

Authors’ Contact Information: Lauren Marietta Chambers, lauren@ischool.berkeley.edu; Diag Davenport, diag@berkeley.edu, University of California, Berkeley, Berkeley, CA, US.

FAccT ’26, Montreal, QC, Canada

ACM ISBN 979-8-4007-2596-8/2026/06
https://doi.org/10.1145/3805689.3806542

such, the students' class project became a critical resource in timely public debates over the city's police budget, and bolstered ongoing demands for greater transparency and accountability into local law enforcement.

***

This episode is instructive not only because it demonstrates a particular learning outcome, but also because it reveals what critical technical practice looks like when it must be enacted rather than merely articulated. The students' work required them to make—and stand behind—decisions under the real constraints and context of policing in Boston: decisions about representation, uncertainty, audience, and timing. This shift, from critique without consequence to decision-making of consequence, forms the starting point for our account of public interest technology clinics as a pedagogy for critical technical practice.

Data-driven work in the public interest hinges not only on critical thinking, but on decision-making: the often invisible series of interpretive, representational, and procedural choices through which data are rendered legible and actionable. Behind each statistical finding presented to a city council, advocacy organization, or public agency lies a dense web of decisions—what the data are understood to represent, how uncertainty is framed, which audiences are centered, and how results might be mobilized by different stakeholders [42, 45, 85]. These choices shape whether technical analyses function as liberatory and productive inputs to public decision-making [17, 26, 40], or instead reinforce existing asymmetries of power, understanding, and control [11, 25, 84].

Since FAccT's founding, substantial attention has been devoted to cultivating critical thinking about technology: critiquing models, surfacing embedded values, identifying harms, and interrogating institutions. Yet the FAccT community has paid far less attention to *how students learn to make decisions* in response to this critical understanding: to act, commit, and take responsibility under uncertainty when technical correctness, ethical clarity, and institutional constraints do not neatly align. As a result, students of critical sociotechnical disciplines may become adept critics of others' ideas while remaining nevertheless underprepared to turn their own ideas into practice.

We argue that **public interest technology (PIT) clinics offer a uniquely powerful pedagogical strategy to develop students into critical technical practitioners.** In a PIT clinic, students undertake semester-long data and tech projects on behalf of public interest 'clients' such as civic organizations or government agencies [20, 40]. Like clinics in other disciplines such as medicine or law, these courses offer supervised practice-based learning [16, 27, 30]. Indeed, our model for public interest technology clinics draws directly from clinical education in public interest law. By situating students in real-world civic contexts where their technical choices have downstream consequences, clinics foreground decision-making as a core component of critical technical practice [38] – a practice in which reflexivity and decisiveness are mutually constitutive [20, 40].

This paper develops a decision-centered account of PIT clinic pedagogy. We propose that what clinics primarily cultivate are not just technical skills [46, 73] or ethical sensibilities [13, 20, 100], but **mental models for decision-making in sociotechnical systems**: ways of reasoning about representation, power, audience, uncertainty, and responsibility that guide action when formal rules or optimal solutions are unavailable. Such mental models are difficult to acquire through abstract instruction alone; they are refined through experience, feedback, and iterative engagement with the consequences of one's own choices—particularly when students are positioned not merely as analysts, but as accountable decision-makers.

Drawing on our collaborative autoethnography of a master's-level PIT clinic at UC Berkeley, and complemented by conversations with instructors at peer institutions, we articulate a pedagogical approach designed to surface and refine these decision-making mental models. We situate this work in prior sociotechnical research (Section 2), detail the mental models central to our teaching philosophy (Section 3), and describe clinic design choices intended to scaffold student decision-making without displacing responsibility (Sections 4–6). At the same time, we emphasize that PIT clinics are not inherently liberatory: poorly designed clinics risk miseducation by reproducing extractive relationships with civic partners, endorsing technological exceptionalism, or obscuring the contested nature of "the public interest" itself. Drawing on Freirean [37], hooksian [43], and Woodsonian [98] perspectives, we analyze how clinic pedagogy can either reinforce or resist these dynamics, depending on how decision authority, accountability, and reflection are structured (Section 7).

Our contribution is threefold: (1) an intervention that centers decision-making—rather than critique alone—as a core pedagogical objective for FAccT-aligned education; (2) a pedagogical theorization linking critical technical practice to the development of decision-making mental models through experiential learning; and (3) a clinic-based case study that surfaces both the promise and risks of teaching civic data practice through real-world engagement. We conclude (Section 8) by inviting the FAccT community to treat decision-making not merely as an object of critique, but as a skill to be taught, practiced, and collectively refined.

2 Related Work

2.1 Critical technical practice

The need for a critique of technology that is enabling rather than disabling—and that moves beyond identifying what is wrong to grappling with what it would mean to *do* something about it—has persisted as a central challenge across decades of sociotechnical research. Scholars have argued and shown that artifacts have politics [3, 94], that algorithmic systems reinforce and augment existing societal disparities [7, 15, 34, 62], and that simplified operationalizations of complex values like fairness are insufficient to meaningfully address sociotechnical problems [45, 49, 59, 76]. FAccT has also begot substantive critiques of itself [5, 9]. Yet such critical insights often leave open a persistent question: how should critique be translated into technical practice that can operate under real-world constraints—where uncertainty cannot be eliminated, stakes are unevenly distributed, and "getting it right" is neither singular nor fully knowable in advance?

This question sits at the heart of Philip Agre's call for a *critical technical practice*, motivated by the claim that the culture and methods of AI research are "actually a powerful force for intellectual conservatism" [2:150] and that any reform must "require a split identity—one foot planted in the craft work of design and the other foot planted in the reflexive work of critique" [2:155]. Subsequent scholarship has elaborated and extended this vision, arguing that the present moment of accelerationist AI calls for practices that resist technosolutionism and instead support collective, context-sensitive forms of action [83]. In response, researchers have developed practice-oriented interventions that embed critique within technical work: approaches that foreground social and community values in design praxes [23, 41, 78, 79, 97]; documentation and reporting artifacts such as model cards [57] and datasheets [39]; and accountability-oriented practices including algorithmic audits and evaluations [51, 69, 88, 92]. Others have articulated roles, actions, and career pathways through which critical technical practitioners can operate within mission-driven and advocacy-oriented contexts [1, 17, 24, 42, 50]. Yet a pedagogical question remains comparatively underdeveloped: how are new sociotechnical practitioners trained to inhabit this split identity? As Shilton argues, values-oriented interventions should "encourage engagement over imagination when at all possible" [79:264]; we contend that this insight applies equally to students.

2.2 FAccT and computing education

Prior FAccT scholars have offered critical analyses of CS, AI, and ML education. For instance, Bilstrup et al. [8] show that AI and ML curricula only rarely surface the design decisions and processes inherent to ML/AI development, while other work documents gaps in how ethics and sociotechnical consequence are taught and assessed [33, 36, 90, 99]. At the same time, scholars have proposed pedagogical innovations intended to improve ethical education in CS [8, 58, 68, 77]. Yet we observe that even these innovations are often tethered to the traditional classroom setting and can struggle to put FATE "concepts, ideas, debates, arguments into practice" [6:430]; indeed, some scholars characterize the current state of CS education as "a funnel from classroom to tech company" [68:516].

The pedagogical philosophy offered in this paper addresses such shortcomings by shifting the learning problem from recognizing sociotechnical critique to practicing decision-making under constraint. From the relative security of a university class, the clinic model allows students to build muscle with strategies that professionals use on-the-job to better understand ethical aspects of data and AI, such as translating prior knowledge and experiences, information foraging, and interpersonal learning [54]. We argue that PIT clinics are the sort of pedagogical model that Raji et al. [68] call for: one that works explicitly with communities, experiential experts, and civic organizations, and that makes the consequences of technical choices difficult to abstract away. Clinics bring the sort of active collaboration and action-oriented partnership with community-oriented and civic organizations that Krafft et al. [51] exemplify into the classroom. In this sense, clinics aim to supplement classroom-based ethical instruction with supervised, consequential practice: an educational setting in which students must learn not only to critique sociotechnical systems, but to act within them.

3 Decision-making under impeachability: proximate and distal mental models

No data product is ever unimpeachable. In civic and public interest data work, additional analysis, refinement, or validation is almost always possible in principle, yet not necessarily warranted in practice. Under real constraints (financial, institutional, political, and more), the central question facing practitioners is not how to eliminate all error, but *when to stop*—and how to recognize when further iteration is unlikely to meaningfully improve decision quality or civic outcomes.

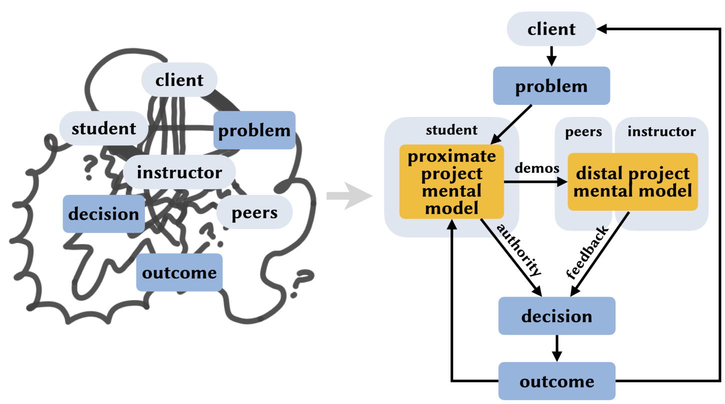

Fig. 1. Idealized development of students' proximate and distal mental models from the start of the clinic (left) to the end (right).

This stopping problem sits at the heart of interdisciplinary economist Herbert Simon's theory of bounded rationality, which emphasizes that real-world decision-makers must act under cognitive, informational, and organizational constraints that preclude full optimization [80, 81]. Simon likens human problem solving to "a search through a vast maze of possibilities" [82:54] and argues that decision-makers must confront a finite question: "When shall I stop the search and accept a solution as satisfactory?" [82:126]. Rather than maximizing, they must *satisfice*: search until a solution meets contextually-defined thresholds of adequacy, then commit. In public interest technology settings, satisficing is not a compromise of rigor but a prerequisite for responsibility: technical excellence divorced from contextual judgment risks delay, misalignment, or irrelevance.

Preparing students for civic and public interest data work therefore requires cultivating decision-oriented *mental models*: simplified internal representations that help practitioners allocate attention, evaluate trade-offs, and decide when additional work is warranted versus when to move forward. Two such models are especially consequential and underdeveloped in conventional technical education: *proximate* and *distal* project mental models. We elaborate on each below, and illustrate their relationships to clinic stakeholders in Figure 1.

3.1 Proximate Project Mental Model: Knowing When to Move On

The *proximate project mental model* governs decision-making when a practitioner is "in the weeds" of a project—responsible for its execution and accountable for its outputs. In this mode, students must distinguish what is non-negotiable from what offers diminishing returns: which uncertainties materially affect downstream use, which modeling choices alter interpretation, and which imperfections are tolerable. Developing this mental model requires learning to impose stopping rules on one's own work. Students must decide when a result is "good enough" relative to the project's goals, audience, and institutional setting. Such judgment is rarely exercised in problem sets, where correctness is academically defined and iteration is implicitly costless. By contrast, client-facing PIT clinics make trade-offs unavoidable: every additional hour spent refining one aspect of an analysis is an hour not spent addressing another, and over-polishing can be as harmful as under-specification when decisions must be made.

3.2 Distal Project Mental Model: Contributing without control

Equally important—but even less emphasized in technical curricula—is the ability to contribute productively to projects one does *not* own. In civic and public interest data ecosystems, practitioners are frequently asked to weigh in on analyses they did not build, with limited time and incomplete context. The *distal project mental model*, which we distill into three elements, governs how to add value in these situations without overstepping, derailing momentum, or offering spurious precision. First, *second-order feedback* prioritizes framing, interpretation, and audience reception over granular technical critique when access to implementation details is limited. Second, *calibration* disciplines the scope and prescriptiveness of feedback to match one's actual proximity to the work, guarding against precise-sounding advice untethered from contextual realities [82]. Third, *active listening* orients feedback toward the project owner's goals and constraints, enabling influence without authority. Without this mental model, practitioners can face a false choice: remain silent unless they can verify every detail, or offer overconfident recommendations that exceed their positionality or expertise.

3.3 Ecosystem Consequences and Pedagogical Implications

Graduates who can only work on their own projects (that is, developing proximate models but not distal ones) hit a ceiling on their individual and collective impact. At the individual level, they are limited to contexts where they retain full control; at the ecosystem level, an overabundance of purely proximate actors means projects are steered exclusively by those already in the weeds, not because outside perspectives are unwelcome, but because potential collaborators lack the mental models needed to contribute effectively at a distance.

These two skills—knowing when to stop iterating on one's own work and knowing how to productively influence work one does not control—are foundational to critical technical practice in the public interest, yet difficult to cultivate through readings, problem sets, or purely hypothetical exercises. Public interest technology clinics are uniquely positioned to develop these mental models by placing students in roles where decision-making is unavoidable and consequences are real. When designed intentionally, clinics do not merely teach students to critique sociotechnical systems, but to act within them: developing judgment about when to push further, when to let go, and how to help others do the same [20, 40].

Consider a common analytic interaction: a student is asked to review a draft visualization prepared by another team for a public-facing report. The chart aggregates program participation by neighborhood, using a color scale that visually exaggerates small differences between adjacent areas. If the student built the chart themselves, the decision is proximate: whether to revise the scale, re-bin the data, add annotations, or accept the imperfection as adequate given time and audience. By contrast, if the student did not build the chart, the decision is distal: how and whether to intervene at all. If they spot a potential misinterpretation, is it material enough to warrant interruption? Should feedback focus on the visual encoding, the accompanying narrative, or the downstream risk of overclaim? And how forcefully should concerns be raised when one lacks control over implementation? Both situations involve judgment under uncertainty, but they demand different mental models. Clinics repeatedly place students in both roles, often within the same project, making the distinction unavoidable rather than abstract.

4 The Public Interest Technology Clinic

The public interest technology clinic adopts a pedagogical style that seeks to teach students how to be critical technological practitioners, going beyond the technical problem-solving offered in traditional technical capstones [46, 73]. Invoking the long history of professional clinical education, in particular in public interest law, a PIT clinic foregrounds professional and real-world skill-building that is distinct from doctrinal, reading-, or discussion-based learning. It further draws from service learning and its commitment to civic education and community service. In this section we compare our approach with other pedagogical models, seeking not to draw arbitrary semantic boundaries but rather to learn from and build upon various traditions.

Definition. We define a public interest technology clinic to be a course that: (1) offers students experiential and project-based education, (2) provides data and technology services to public interest organizations, and (3) aims to develop students' skills and mental models as PIT practitioners. We embrace a definition of the 'public interest' that is actively negotiated and contested [89]. This definition encapsulates not only our own clinic at Berkeley but also similar endeavors at other institutions,[2] for instance, the Olin College "PInT" initiative [20, 40] or the BU Spark! Lab at Boston University.[3]

4.1 Public Interest Technology

The specific term 'public interest technology' has grown in popularity in large part due to philanthropic investment by major players including the Ford and MacArthur Foundations in 2015, as well as the foundation of the Public Interest Technology University Network (PIT-UN) in 2019. PIT is intentionally designed to be a 'big tent' that includes individuals across domains, sectors, and expertise. It is no coincidence that the popularity of PIT has grown alongside the magnitude of the so-called 'techlash' [91] wherein public awareness has grown around worker mistreatment [55, 70], extractive data practices [101], contracting with militaries [31], and alignment with political conservatism among Big Tech companies including Facebook (now Meta), Google, Amazon and others [52]. PIT is but part of a broader moment that includes the development of interdisciplinary communities of practice in critical technology studies (e.g., the first FAccT conference in 2018, the Data & Society research institute established in 2014) and the establishment of new nonprofits that are both tech-minded and advocacy-oriented (e.g., Code for America in 2009 [75], the Algorithmic Justice League in 2016 [14], Data for Black Lives in 2017 [56], etc…). The PIT clinic applies a clinical model to technical work in direct response to this growing momentum towards, and demand for, initiatives and opportunities that design and deploy tech in the public interest.

4.2 Similar models

4.2.1 Profession-based clinical education. Clinics are a common pedagogical structure in professional education across disciplines, including in the instruction of law, nursing, psychology, physiotherapy, and more [16, 27]. Of course, the approaches and elements of an environmental law clinic will be distinct from those of a speech pathology clinic. Yet one element that is nonetheless shared across disciplinary clinics is, as Cantatore et al. describe, "the centrality of practice-based learning" that results from "immersion into the profession" under instructor supervision [16:5].

The nearest methodological neighbor of the PIT clinic is the cybersecurity clinic [87]. Cybersecurity clinics "bridg[e] the gap between academic learning and real-world public service, aligning student education with community-based cybersecurity interventions that enhance the resilience of public-serving organizations" [61:244]. A formal international Consortium of Cybersecurity Clinics[4] was founded in 2021. As of early 2026,

[2] The first author maintains a database of known public interest technology clinics here: http://tiny.cc/pit-clinic-database

[3] See "Our Courses" at the BU Spark! Lab, October 3, 2025.

[4] https://cybersecurityclinics.org/

the Consortium includes 60 active members that offer cybersecurity clinics in their institutions [61]. These clinics, like PIT clinics, employ data- and tech-based methods as their tools of choice as they address clients' problems.

The political and societal alignments of the PIT clinic, however, inherit explicitly from the public interest law clinic [10, 30, 48, 95]. Clinical legal education took off in the United States in the 1960s and 1970s at the peak of the civil rights and anti-war movements. Students demanded a more radical form of legal education that was responsive to the plain inequities of the time, and the Ford Foundation stepped up to fund 'socially progressive' clinical education [30, 95]. As the years passed, not all law school clinics stayed close to these progressive roots; on the whole clinics became "less ideological and more pragmatic" [30:1467] in the 1980s as Ford funding ran out and in response to American Bar Association critiques that "law schools overemphasize doctrine at the expense of skills and values" [30:1468]. Nevertheless, many law schools today operate multiple clinics, some with explicit social justice aims such as environmental justice, criminal justice, human rights, or housing justice [10, 30, 48]. Clinical law professor Jon Dubin articulates three goals of what we might call public interest law clinics: (1) to provide services or pursue reform on behalf of underserved clients and communities, (2) to expose students to a public service ethos, and (3) to develop students' personal sense of justice through experiencing "the impact of the legal system on subordinated persons and groups" [30:1475-7]. These goals mirror the three criteria we attribute above to PIT clinics. The service goals of public interest law clinics are also educational goals [95], an orientation that Dubin observes [30:1481] is shared by another pedagogical model: service learning.

4.2.2 Service learning. Service learning is a pedagogical approach which "combines academic study with community service" [35:517]. It became a popular pedagogical approach beginning in the 1990s as an effort to "contribute to engaged citizenship" [35:519]. While it is, at minimum, an enterprise in which both students and community partners benefit, at its best it has the promise to be an experience which is truly "mutually transformative" [13:108]. Research suggests that service learning has the potential to effectively train graduates with not only relevant skills and knowledge, but also an understanding of social issues and a disposition for public interest careers [13, 100]. Service learning can have particularly meaningful impacts on more junior students, such as first-years [12, 22].

Service learning employed within the specific environments of computer and information science (CIS) education has been the focus of specific study by various Computer Science Education scholars. Robledo Yamamoto et al.s' 2023 meta-analysis of over 80 such classes identifies four common goals expressed in CIS service learning courses including: (1) enhancing student learning with hands-on, real-life experiences, (2) addressing community needs, (3) developing students' civic responsibility, and (4) building mutually-beneficial relationships with community partners [71:7]. Studies have also shown that service learning can be an effective means of recruiting underrepresented students to computer science classes that otherwise frequently struggle with student diversity [63, 66].

5 Methodology

This paper draws on collaborative instructor autoethnography [18] to examine the pedagogical design and implementation of a master's-level Public Interest Technology (PIT) clinic. Autoethnography is a qualitative method that "seeks to describe and systematically analyze personal experience in order to understand cultural experience" [32] and that has been adopted in prior pedagogical research [6, 29, 86]. In this case, the "culture" under examination is the classroom culture that emerged from the implementation of the Berkeley Clinic. An ethnographic approach combines sustained instructor reflection with peer dialogue and artifact analysis to surface recurring pedagogical patterns, decision points, and tensions that would be otherwise difficult to capture. For discussion of the authors' positionalities, see Section 9.1.

5.1 Data Sources and Analytic Process

Our analysis draws on four primary sources of evidence. First, *reflective documentation*: over 18 months of project formation and two semesters of student instruction, both authors maintained contemporaneous field notes documenting pedagogical decisions, student interactions, moments of uncertainty or conflict, and project outcomes. These notes were revisited iteratively during manuscript preparation to identify recurring themes and critical incidents. Second, *peer dialogue*: to mitigate the limitations of our Berkeley context, we engaged in informal but substantive conversations with instructors running analogous PIT or clinic-style courses at peer institutions. These discussions were used to validate observations, surface alternative design choices, and identify patterns that transcend local context. Third, *artifact analysis*: we analyzed course materials (including syllabi, discussion prompts, student presentations), student deliverables, and partner communications that resulted from the operation of the course. These artifacts offer insight into decision-making processes and pedagogical or institutional trade-offs. Fourth, *retrospective synthesis*: during manuscript development, the authors collaboratively identified recurring pedagogical tensions, critical moments, and decision-making patterns across semesters.

5.2 Limitations

As instructor-researchers (see §9.1), we occupy a position of structural power that shapes what we observe and how we interpret it. We do not claim access to students' internal experiences or learning processes beyond what is visible through their actions, artifacts, and interactions. Our account therefore privileges instructor perspective and may miss dynamics legible only to students or community partners - a shortcoming across critical CS education research [72]. Additionally, autoethnography prioritizes depth, reflexivity, and analytic transparency over generalizability [18]. Our findings emerge from a specific institutional and disciplinary context and may not transfer straightforwardly to other settings without adaptation.

6 The Case of the UC Berkeley PIT Clinic

6.1 Setting and Structure

The UC Berkeley Public Interest Technology clinic[5] is offered as a three-credit course by the School of Information and the Goldman School of Public Policy in which student teams partner with public-sector and nonprofit organizations on semester-long, data-driven projects. Its approach is explicitly interdisciplinary [68]; as such, it enrolls graduate students from diverse disciplinary backgrounds including data science, information science, public policy, and related fields. Projects span domains such as public benefits access, civic outreach, legal services, and program evaluation, and typically involve substantial uncertainty around data availability, stakeholder priorities, and institutional constraints. The clinic began in Spring 2025 and is held every semester, allowing certain projects to continue across semesters, or certain students to participate multiple times with a cascading mentorship model [47]. Over the past three semesters and counting, we incorporate student feedback to improve each iteration of the clinic.

The Berkeley clinic is structured around a deliberate separation between pre-semester project formation and in-semester project execution, with the instructor's role changing fundamentally across these phases. This separation is central to how our clinic scaffolds students' development of proximate and distal decision-making mental models. We note this is a distinct pedagogical choice from other PIT clinics; Olin's PInT initiative puts students fully in charge of their own project placement, requiring them to identify and coordinate with prospective community partner clients at the start of the semester [40].

We also note that a detailed description of more specific implementational logistics - e.g., regarding class size, client selection, problem formulation, project scoping - is out of scope for this higher-level paper. We hope future work might offer more detailed exploration of such nuts-and-bolts elements.

6.2 Instructor Role Before the Semester: Proximate Judgment in Project Formation

Before the semester begins, the instructor takes primary responsibility for sourcing, scoping, and setting up potential projects. This involves extensive outreach to government offices and nonprofit organizations, particularly

[5]For more information, see the clinic website: http://tiny.cc/berkeley-pit-clinic

in adjacent communities, and centers listening rather than imposing solutions. Initial conversations often begin with a deliberately open prompt—"If you had a team of highly skilled graduate students working on the most impactful problem for you, what would that be?"—with the expectation that first answers are rarely immediately actionable. Many prospective partners face structural constraints that prevent them from easily articulating technical problem statements: limited staff capacity, lack of analytic resources, or uncertainty about what data could realistically support. The instructor's task at this stage is therefore not to define the project unilaterally, but to iteratively refine an initial intuition into a problem that is both substantively meaningful and pedagogically tractable. This includes judging whether a project is sufficiently scoped, whether meaningful progress can be made within a semester, and whether the ambiguity involved aligns with the clinic's learning objectives.

In this pre-semester phase, the instructor necessarily operates in a proximate decision-making mode (see §3.1): exercising judgment about trade-offs, feasibility, and stopping points on behalf of future student teams. The goal is not to eliminate uncertainty, but to ensure that what remains is productive uncertainty that students can engage with, rather than ambiguity that stymies progress.

6.3 Instructor Role During the Semester: Modeling Distal Judgment and Ceding Control

Once the semester begins and students are onboarded to projects, the instructor's role shifts sharply. Decision authority over technical direction, workflow, and client engagement is transferred to student teams wherever possible. The instructor adopts a deliberately low-touch posture, resisting the impulse to dictate solutions or intervene prematurely, even when students struggle.

Biweekly class sessions center on short project demos, typically structured as five minutes of presentation followed by ten minutes of questioning and feedback from other students. Presenters are required not only to share their outputs, but to defend their decisions: why a particular assumption was made, why a modeling path was pursued or abandoned, why a visualization was designed in a given way. These sessions are frequently uncomfortable. Students accustomed to well-specified assignments often freeze when confronted with open-ended responsibility—for organizing files, coordinating across skill sets, deciding when to contact a client, or determining when a first pass is "good enough" to share. Yet our philosophy here aligns with that of the Olin clinic: "the value of PInT's clinic program is actually in students struggling with those questions. This is what public interest technology and engineering as public work looks like. It is messy" [40:86].

A central instructional task at this stage is emotional regulation in service of decision-making. Early confusion and frustration are treated not as failure, but as predictable responses to new, consequential authority. Rather than resolving uncertainty for students, the instructor intervenes by asking questions that return decisional power to them. Over time, students become increasingly invested in their work, exhibiting what might be described as an "IKEA effect" [60]: ownership emerges precisely because they are the ones making the calls.

Crucially, during this phase the instructor models a distal mental model (see §3.2), offering feedback on framing, interpretation, and downstream implications rather than prescriptive or granular technical fixes. The instructor's authority is exercised not through control but through calibration: signaling concerns, highlighting trade-offs, and inviting reflection.

6.4 Relational Norms and the Public Interest

From the outset, the clinic emphasizes relationship-building as a core component of public interest work. A recurring refrain—"we are people helping people help people"—anchors classroom culture. While producing something useful for client organizations is an important intermediate goal, the clinic's north star is the recognition that data work matters only insofar as it affects real people's lives. This is a philosophy shared across service learning initiatives in computer and information science: "ICT by itself will not improve or increase democracy, equality, social inclusion, or any other social good" [21]. This framing shapes how students interact with one another, with clients, and with feedback. Students are encouraged to share incomplete work early ("share bad charts"), to treat critique as a property of their artifacts rather than themselves, and to communicate uncertainty without defensiveness. These norms are reinforced through mid-semester reflection sessions explicitly dedicated to discussing and reflecting on students' challenges, emotions, and decision-making struggles.

6.5 Design Choices in Practice: Learning to Decide by Deciding

Across projects, students repeatedly confront the same uncomfortable realization: there is no unimpeachable outcome, only decisions that must be made under uncertainty. As one student reflected during a collective class debrief, "It feels like I've been looking at pictures of horses my whole life, and this class is the first time I've actually gotten to see and interact with one." Here we share specific vignettes (★) that exemplify how consequential decisions facilitate learning.

★ *Operationalizing social constructs.* In one project, a housing nonprofit shared a decade-long dataset detailing millions of dollars in home renovation spending, and requested that students help assess the organization's success in achieving their primary mission: reducing disparities in intergenerational wealth. Yet the data contained no direct measures of household wealth. Students immediately faced a dilemma that no problem set could resolve for them: how much inference is too much? Should they attempt to estimate wealth using proxies such as property values, tax assessments, or third-party sources like Zillow? What biases would those introduce? Was a noisy, partial answer preferable to nothing—and if so, under what framing? Rather than supplying a correct path forward, the instructor "wondered together" with the students, surfacing questions they would need to answer themselves: What is the standard definition of wealth? What does it omit? Which components are most observable and most correlated with the underlying construct? What errors would be introduced by measurement bias, and how would those errors propagate into recommendations? Students eventually learned that choosing a numerical proxy is itself a decision with ethical and political consequences, not merely a technical workaround.

★ *Iterative problem formulation.* In another project, a team building a Qualtrics-based experimental testbed for detecting algorithmic bias in healthcare encountered a persistent implementation bug. During a demo, peers began asking clarifying questions (repeated cycles of explanation, critique, and reformulation) that revealed a deeper issue: the "bug" was not in the code but in how the task had been conceptualized. Early in the semester, such moments were incapacitating; by the final weeks, they had become routine. Students are given the opportunity to treat confusion not as failure, but as diagnostic signal—evidence that a decision boundary needed to be redrawn.

★ *(Literally) Bounding analysis.* A third team, conducting a cost-benefit analysis for an ML system to prioritize outreach for tax refund uptake, faced a subtler decision. After running sensitivity analyses, they were asked about the x-axis range in a key figure. A wider range would show robustness across a broader space of assumptions—but at what point did additional range cease to be meaningful? How much deliberation was warranted before simply "calling it" and moving on? This led to a meta-discussion about when further thinking improves judgment and when it merely delays commitment—a distinction students reported never having to make explicit before.

These decisions were scaffolded by a weekly stand-up structure designed to surface students' own mental models. Each week, each team member reported: (i) what they had done, (ii) what they were working on and how, and (iii) whether they considered themselves blocked or progressing. Crucially, "blocked" was treated not as a failure state but as a signal that a problem lay outside the student's scope of control. Instructor intervention focused first on reframing: Is there an alternative path that allows progress despite the blockage? What would it cost—in time, credibility, or opportunity—to wait? What preparatory work could be done now? Two teams experienced this in starkly different ways.

★ *Making progress despite setbacks.* One team working on a tax policy project was blocked for nearly two months by the longest U.S. government shutdown in history. Unable to access data or even communicate with their client, students initially believed progress was impossible. Instead, they constructed a synthetic dataset

Table 1. Six Design Principles for Public Interest Technology Clinics

Principle	Pedagogical Goal	Clinic Design
Productive Uncertainty	Preserve ambiguity that students can engage with, rather than eliminating uncertainty entirely or introducing ambiguity so severe that it forecloses action.	Projects are scoped to ensure that uncertainty, surprises [29], and "disorienting moments" [30:1478] are substantive—about framing, trade-offs, or interpretation—yet bounded enough that students can still make progress and experience the consequences of their decisions.
Transferred Authority	Transform students' ability to critique into a practice of judgment and action.	Once the semester begins, while the instructor remains accountable for the overall learning environment, students are positioned as the primary decision-making agents responsible for technical direction, workflow, and client engagement.
Consequential Choices	Students make complex decisions of consequence, motivating their sense of responsibility, rather than worrying about making the 'correct' decisions.	Student decisions have real downstream effects on actual partner organizations. The clinic resists simulated or purely hypothetical projects in favor of engagements where choices around scope, framing, and timing matter beyond the classroom.
Structured Reflection	Critical reflection [28] asks students to build comfort articulating not only what they did, but why and how they evaluated trade-offs under constraint.	Regular demos, stand-ups, and reflective discussions are embedded into the weekly rhythm of the course rather than relegated to post hoc evaluation. End-of-semester grades hinge on student self-evaluations regarding the development of their mental models.
Modeling Distal Practice	Students observe how calibrated, second-order engagement can add value even when they lack direct ownership of a project.	The instructor models how to contribute to projects without controlling them. Feedback for students emphasizes framing, audience, interpretation, and downstream implications rather than prescriptive technical fixes.
Relational Accountability	Students foreground respect for clients' time, humility about technical expertise, and accountability to affected publics.	The clinic frames its work as fundamentally relational—"people helping people help people"—rather than as technical service delivery. Grading schemes reflect this, rather than reinforcing outcome-based notions of in-class success or failure.

Icons from fontawesome.com

and used it to design and iterate on every planned analysis and visualization. They mapped out post-shutdown communication strategies, weighed the ethics of rushing a client for end-of-semester deadlines that mattered only to students, and prepared code to run immediately once access resumed. When the government reopened, the client—impressed by the forethought and care—prioritized the collaboration, enabling the team to deliver a complete analysis by semester's end.

★ *Deciding whether to give up.* By contrast, another team, frustrated by prolonged data delays from a public safety agency, disbanded and withdrew from the course. Importantly, this was not treated as a failed pedagogical outcome. Students who left the course encountered early—and viscerally—the realities of data agreements, bureaucratic delay, and institutional friction. Those who stayed witnessed how much progress could still be made under constraint, and were forced to ask themselves whether such work was worth pursuing despite its frustrations. Graeff and Wood [40] describe a complementary example from the Olin clinic: at the end of one semester's project, students ultimately *refused* to build a sex trafficking identification tool that they had originally set out to build for fear it might endanger sex workers - and this outcome was celebrated by the clinic's faculty advisor [20, 40]. In the clinic environment, there can be a lot to learn from giving up.

Many students enter public interest work with deeply romanticized expectations. The clinic repeatedly surfaced the difference between what one instructor analogized as a romantic crush versus a committed relationship: the former sustained by idealized projections, the latter grounded in a clear-eyed assessment of imperfections and trade-offs. Experiencing that distinction—rather than merely being told about it—proved formative. As one student noted in an end-of-semester reflection, the course taught them to "get work done in ambiguous circumstances" and to "operate in sub-optimal, imperfect contexts," even when progress felt uncertain.

Across these examples, *students did not learn effective critical decision-making by reading about it.* They learned by repeatedly encountering situations where no one could tell them the right answer, or where moving forward required committing to an imperfect choice and defending it to peers, clients, and themselves. Over time, the uncertainty and stagnation that characterized students' early weeks gave way to fluency: students became more concise in stand-ups, more calibrated in feedback, and more resilient in the face of uncertainty. By the end of the semester, they expressed pride in what they completed in their projects - however imperfect. This is what critical technical practice looks like when it is lived rather than described.

6.6 Design Principles for Public Interest Technology Clinics

Taken together, the preceding cases illustrate not a set of isolated instructional tactics, but a coherent pedagogical logic. We distill from this experience a small set of recurring principles that capture how the clinic supports students' development of decision-making mental models. These principles, listed in Table 1, are not intended to be a prescriptive checklist but as descriptions of the conditions under which clinic-based learning consistently became productive rather than stagnant, extractive, or superficial. These six principles help explain how PIT clinics can cultivate decision-making capacity rather than mere technical or critical proficiency. They also clarify why clinics that lack one or more of these features risk miseducation: reproducing technocratic authority, insulating students from consequence, or treating public interest work as an abstract exercise rather than a lived civic practice.

7 Discussion: From Pedagogy to Practice

Sections 3–6 argued that PIT clinics are uniquely positioned to cultivate decision-making mental models: ways of reasoning about uncertainty, representation, audience, and responsibility when no option is unimpeachable. The ultimate goal is to cultivate, in Agre's words, critical technical practitioners [2]: multidisciplinary experts who can move through the world as not only adept data-crunchers but also as thoughtful and effective problem-solvers. In this section, we clarify the good, the hard, and the ugly of fostering critical technical practice via this clinical model. Informed by critical educational theorists as well as critical data scholars, we articulate the promise of liberatory outcomes, challenges of avoiding miseducation, and the threat of hegemonic power dynamics.

7.1 The Good: Potential for Liberatory Pedagogy

Black feminist scholar bell hooks describes education as "the practice of freedom" [43:152]. Under this lens, what we view to be potentially 'liberatory' about a PIT clinic is not that it guarantees virtuous material outcomes or resolves structural inequities in one semester – but that it reorganizes how students relate to knowledge and power.

7.1.1 Developing a personal sense of (data) justice. In law school, clinical instructors are clear that a primary goal of public interest law clinics is for students to experience, up close, how broken the legal system is and to be disoriented or disturbed by what they witness. The point is, in part, to undergo a resulting "perspective transformation" in which they "engage[] in critical thinking focusing on reassessment of societal and personal beliefs, values and norms" [30, 67:52], and ultimately to develop a more robust "personal sense of justice" [30:1477]. In the PIT clinic, too, the discomfort is the point. Negative feelings, in data science as in other contexts, can serve as meaningful and instructive signals worth attending to [53]. The on-the-ground work in PIT clinics seeks to challenge students' *a priori* assumptions about what data and technical solutions are good for, how effective they can be, and how difficult it is to meaningfully address societal problems. Via *relational accountability* that centers community and *consequential choices* that mean the stakes are real, students ideally come to realize how datasets are often incomplete or unhelpful, how quantitative methods can abstract away nuance to the point of uselessness, and how every choice in the design of sociotechnical systems can have ethical and political implications. In this way, clinics can be a meaningful response to Brazilian educator and Marxist philosopher Paolo Friere's famed rejection of the so-called banking model of education, in which "knowledge is a gift bestowed by those who consider themselves knowledgeable upon those whom they consider to know nothing" [37:58]. Friere instead champions a participatory, dialogic learning that develops students' *conscientizaçao* or 'conscientization' that deepens students' "attitude of awareness" [37:109] - very much like, we argue, what a public interest clinic seeks to achieve.

7.1.2 Critical reflection. Education scholars are clear on the value of intentional reflexivity for students. hooks quotes Friere on this matter: "praxis is not blind action, deprived of intention or of finality. It is action and reflection" [43:48]. Clinical educators Delany and Molloy define critical reflection as an active skill that hinges on "the process of thinking about practice" and results in deeper understanding and more sophisticated critique of "established practice knowledge" [28:4]. This reflection can take a variety of forms in classrooms, though Delany and Molloy [28] synthesize two main methods of integrating student reflection into healthcare clinics: reflective diaries/journals and verbal feedback between students and instructors. When PIT clinics require *structured reflection* via in-class demos, critical discussions, and end-of semester self-evaluations. students develop stronger mental models for critical technical practice: articulating not just what they have done, but why, and what resulted.

7.1.3 Training future gatekeepers. Finally, PIT clinics should not be understood as serving only students who plan to enter explicitly civil- or public-sector roles. Many students studying computing will move into industry, including large technology companies. If anything, however, this makes the development of critical technical practice *more* urgent, not less. Industry-bound graduates often gain the power to shape - at scale - what problems are prioritized, which metrics define success, which stakeholders are heard, and which harms are treated as 'externalities' [93, 96]. In other words, they become gatekeepers of sociotechnical reality. Legal clinics, again, offer a useful analogy. Clinical courses are widely understood to make students better lawyers overall, not only better public interest lawyers, because they train professional judgment: how to act under constraint, how to be accountable to clients, and how to navigate uncertainty without surrendering responsibility [19]. Similarly, PIT clinics can be understood as one means of professional formation for the broader sociotechnical ecosystem, regardless of sector. Asking future Big Tech employees to practice *intentional*, *consequential*, and *accountable* sociotechnical decision-making helps to train not only better workers, but better citizens.

7.2 The Hard: Miseducation Risks

In his 1933 book, *The Miseducation of the Negro* [98],[6] Black historian Carter Woodson argues that the prevailing educational system in the early 20th century was deliberately structured to train Black Americans to occupy a subservient place in the social order, rather than to think critically and advance of their communities [98:6-8]. Woodson asserted that education must go beyond "the mere imparting of information" and ultimately seek to teach a student to "think and do for [them]self" [98:3]. Here we consider challenges in the PIT clinic model that risk producing graduates who emerge technically trained yet politically unmoored, disempowered, and disconnected from the economic, political, and social consequences of their work.

7.2.1 Navigating difficult emotions. Students often encounter that *productive uncertainty* doesn't feel good.[7] Yet while protecting students from uncertainties or offering them more abstracted projects might reduce student discomfort, it also reduces the relevance of the pedagogical approach to real praxis. Indeed, discomfort and disorientation can be generative emotions in clinic settings (see Section 7.1.1). Apathy and boredom, however,

[6]This work inspired the title of Lauryn Hill's septuple-platinum 1998 album.

[7]Perhaps curiously, we share that, at least in our Berkeley clinic, very real student struggles and discomfort have not translated to poor course evaluations. However, we acknowledge that concerns about negative evaluations might prompt valid hesitation about the model we describe here, particularly for adjunct or pre-tenure faculty who experience professional precarity.

are more dangerous. In PIT clinics, this can look like stagnation, a feeling of not advancing in projects, or students being too in over their head. hooks argues that it is a hallmark of liberatory pedagogy that instructors take responsibility for boredom and disinterest among students, and adjust curricula accordingly [43:155-6]. Instructors of PIT clinics must seek to strike a tricky balance: between offering helpful *distal feedback* to help students get and remain un-stuck, and maintaining students' *authority* over their projects and decisions.

7.2.2 Avoiding instructor dependence. Relatedly, STEM students are often used to their instructors telling them which problems to solve, and whether or not they've gotten an answer 'right'. *Transferred authority*, in which students pull the strings and call the shots in their projects, and *modeling distal practice*, in which instructors are but backseat drivers, are not the norm in classroom settings. The client-facing nature of clinics might tempt instructors to take a higher-touch approach managing student projects, yet we and other PIT clinicians argue it is imperative that instructors resist this urge [20, 40]. Perhaps moreso than in other lines of computing work, public interest tech practitioners often work on small or even one-person teams with limited peer support or mentorship [17]; PIT clinics can prepare them to be independent practitioners. As Olin professors Graeff and Wood say: "expect a lot from your students. Show that you believe in their ability to tackle complex issues and work through ambiguity. They will rise to the occasion and grow through the experience" [40:92-93].

7.2.3 Lacking a true profession. Legal and medical clinics teach students norms of praxis as decided and shaped by a central examination process, licensing body, and professional association. Of course, no such corollaries exist in computing or data science. And while some practitioners are making an effort to professionalize public interest technology specifically, such initiatives are still green and often under-resourced [17, 50]. As a result, there are few (or, arguably, no) guardrails or agreed-upon standards regarding ethical and professional requirements for data scientists, software developers, and the like – and no roadmap to guide PIT clinic priorities. Relatedly, various scholars have observed that well-intentioned efforts to teach public interest technology topics can be hampered by would-be instructors' limited familiarity with critical sociopolitical analysis or even simple project management, due to mono-disciplinary training and inexperience outside of academia [6, 68]. These ambiguities underscore the importance of students' experience with *productive uncertainty* and *transferred authority* to develop strong mental models in clinics: they will need them in the wild west of professional public interest tech work.

7.2.4 The cost of clinics. The novelty of client-facing tech clinics poses material challenges as well. Experiential, partner-based teaching is structurally more resource-intensive than simulated projects and requires: lower student-to-faculty ratios; sustained relationship-building; and institutional support for scoping, coordination, and care. This work can be substantial. Ideally, departments would assign dedicated staff to carefully execute the logistical elements of partnership management rather than leaving such tasks to oversubscribed faculty or RAs. Yet, particularly as many U.S. universities face budget cutbacks in 2026, PIT clinics must often justify why 'making up projects' is not equivalent. The answer, as this paper demonstrates, is that critical technical practice cannot be simulated cheaply: effective pedagogy depends on real *consequence*, real *uncertainty*, and real *relational* stakes. Yet administrators might require convincing.

7.3 The Ugly: Exceptionalism and Extractionism

Service learning scholars Bringle and Clayton describe aspiring to *thick reciprocity* in which "all teach, and all learn, all serve, and all are served" [13:108]. When clinics are designed for thick reciprocity, they can expand partner organizations' capacity to mobilize data, communicate evidence, and navigate institutional constraints [17, 26, 40]. But if partner organizations function primarily as thin props for student learning, clinics risk reproducing extractive relationships between elite universities and under-resourced communities. While, ideally, the creation and operation of PIT clinics constitutes a form of community building in and of itself, reality can be more complicated. Without *relational accountability*, they risk real costs: partner time lost, trust eroded, and capacity depleted [11, 25, 84].

7.3.1 Hierarchies of knowledge and institutions. PIT clinics sit within universities that often operate with formidable resources, credibility, and reputational power. That positioning enables clinics to broker relationships and mobilize student labor – but it also raises questions about extraction, access, and who gets to define the 'public interest' [84]. As Irani and Silberman narrate in their account of designing the Turkopticon project, dominant narratives about what forms of labor are more and less valuable or what sorts of institutions confer credibility are difficult to unweave [44]. Efforts for clinics to remain *accountable* to clients go some way toward this tension. Yet even more fundamentally, the language of data and computing is already mired in epistemological hierarchies. Privileging what Crooks and Currie call 'datalogical' arguments can displace and eclipse other forms of (especially lived, experiential, and narrative) knowledge [25, 64, 65]. This is why it is so crucial that clinics integrate critical *student reflexivity* – reflecting upon not just the work of the clinic but also its limitations and political valences. Even so, these are incomplete fixes for deeply entrenched orientations of power.

7.3.2 Tech exceptionalism and 'unicorns'. In 2026, tech companies have captured not only global markets but also societal imaginations [15]. Many theorists and laypeople alike subscribe to a dangerous exceptionalist view of scientistic information as fundamentally distinct from culture [74]. Relatedly, Raji et al. [68] criticize the individualistic trope of the 'ethics unicorn,' the lone sociotechnical expert who is capable of solving any tech ethics problem. Such a unicorn is, of course, a myth; yet if clinics tacitly portray technical intervention as the central driver of civic progress without critical *reflection* or *relational accountability*, they can reinforce the myth of the savior technologist or the imaginaries of omnipotent technology. Such mythos obscures the reality that 'public interest' is contested [89] and that civic outcomes are shaped by institutions, politics, and collective action more than by tools.

8 Conclusion and Call to Action: Teaching Decision-Making, Not Just Critique

This paper argued that FAccT-aligned education has made advances in cultivating critical thinking about sociotechnical systems, but has paid less attention to cultivating students' *decision-making* as an element of *critical technical practice*: the capacity to act, commit, and remain accountable when no technical solution is unimpeachable.

Via collaborative autoethnography of a master's-level PIT clinic, we develop a decision-centered account of a clinic, and propose that such pedagogies can cultivate decision-making mental models that are central to actualizing critical technical practice – if difficult to acquire through conventional coursework. We then turn to educational theorists including Freire, Woodson, and hooks to help understand how clinics can become more or less liberatory, miseducative, or extractive depending on how student authority, uncertainty, reflection, and accountability are structured. We celebrate the liberation that can come from student *conscientizaçao* and reflexivity, while cautioning against coddling or micro-managing students, and warning about the dangers of reinforcing entrenched hierarchies around universities and technology.

We conclude by inviting the FAccT community to **treat decision-making not merely as an object of critique, but as a skill to be taught, practiced, and collectively refined**. This includes innovating pedagogical designs (see Table 1) that operationalize accountability, develop norms for non-extractive partnership, and build institutional arguments for why clinics—despite their costs—are a worthwhile infrastructure for responsible professional development. Ultimately, the value of PIT clinics lies not in resolving frictions between higher education, public interest, and civic tech, but in intentionally navigating them—and equipping students to do the same, especially as many will carry their judgment into positions of real power across the sociotechnical ecosystem.

9 Endmatter

9.1 Positionality statements

The first author is a senior graduate student in information science. She played a supportive role in the clinic, holding office hours for students, attending planning meetings with the primary instructor, offering a data visualization workshop each semester, and (in the second semester) helping to facilitate weekly student project demos. She has professional experience working as a software developer and data analyst in both scientific and nonprofit political advocacy contexts. Her graduate research focuses specifically on the practice of public interest tech, and informs a pedagogical commitment to preparing students for the reality of professional work in PIT.

The second author and primary instructor is a faculty member trained in behavioral economics, data science, and public policy, with research focused on sociotechnical decision-making, algorithmic governance, and public-interest data systems. Alongside academic work, he has collaborated extensively with governments, nonprofits, and advocacy organizations on applied data projects related to public benefits access, legal services, and civic infrastructure.

The pedagogical commitments articulated in this paper—including ceding decision authority to students, foregrounding uncertainty, and emphasizing accountability to affected publics—are informed by an explicit effort by both authors to counter technocratic default assumptions and to approach public interest technology as a relational, value-laden practice rather than a neutral technical service. Both authors also operate from positions of relative institutional security and authority within a large public research university—an asymmetry that both enables access to resources and civic partners and shapes power relations among instructor, students, and community organizations.

9.2 Ethical considerations

As this project constitutes a collaborative instructor autoethnography, we did not require student consent for involvement in the clinic, and accordingly restrict our analysis to instructor decision-making, pedagogical design choices, and observable project dynamics. Any data collected about students by the instructor-researchers did not go beyond typical data collection for the design, operation, and evaluation of a university course. Relatedly, this paper does not report student learning outcomes or assess individual performance. Where student experiences are discussed, they are described at a level of aggregation and / or abstraction sufficient to protect individual privacy while still illuminating pedagogical phenomena.

9.3 Generative AI usage statement

Generative AI tools were used solely to assist with grammar and fluency of writing, including copy editing for clarity and concision. Generative AI tools were not used to generate substantive text or intellectual content. All research questions, arguments, analysis, interpretation, and conclusions were developed by the authors, who reviewed and take full responsibility for the final manuscript.

Acknowledgments

Conversations with and feedback from many people helped shape this paper including: Safi Aharoni, Frank Abissi, Elijah Baucom, Fran Berman, Henry Brady, David Coleman, Ziba Cranmer, Adam Finn, Laurel Fletcher, Eric Gordon, Erhardt Graeff, Michael Chang, Peter Linquiti, Deirdre Mulligan, Michael O'Hare, Felix Owusu, Birin Padam, Manuel Sabin, Portia Taylor, Valentina Rozo, Marisella Rodriguez, AFOG participants, Berkman Klein Center Public Interest Technology Working Group (BKCPITWG), and the anonymous reviewers.

This material is based upon work supported by the first author's affiliation with the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 2146752. Any opinions, findings, and 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. The first author was further supported by the Pre-Doctoral Fellowship from the Ford Foundation (2023-28) and the UC Berkeley American Cultures Engaged Scholarship (ACES) Chancellor's Fellowship (Fall 2025).

References

[1] Rediet Abebe, Solon Barocas, Jon Kleinberg, Karen Levy, Manish Raghavan, and David G. Robinson. 2020. Roles for computing in social change. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 252–260. https://doi.org/10.1145/3351095.3372871

[2] Philip E. Agre. 1997. Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI. In *Social Science, Technical Systems, and Cooperative Work: Beyond the Great Divide*, Geoffrey C. Bowker, Leslie Gasser, Susan Leigh Star, and William Turner (Eds.). Lawrence Erlbaum Associates.

[3] Madeleine Akrich. 1992. The De-Scription of Technical Objects. In *Shaping Technology / Building Society: Studies in Sociotechnical Change*, Wiebe E. Bijker, John Law, Trevor Pinch, and Rebecca Slayton (Eds.). MIT Press, Cambridge, MA, USA, 208.

[4] American Civil Liberties Union of Massachusetts, Progressive Massachusetts, Councilor Ricardo Arroyo, and Boston University Spark! Lab. 2020. *Boston Police Department Court Overtime*. Technical Report. https://data.aclum.org/wp-content/uploads/2020/11/BPD-Court-Overtime-Report-ACLUM-Progressive-Mass-BU.pdf

[5] Chelsea Barabas, Colin Doyle, Jb Rubinovitz, and Karthik Dinakar. 2020. Studying up: reorienting the study of algorithmic fairness around issues of power. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. ACM, Barcelona Spain, 167–176. https://doi.org/10.1145/3351095.3372859

[6] Jo Bates, David Cameron, Alessandro Checco, Paul Clough, Frank Hopfgartner, Suvodeep Mazumdar, Laura Sbaffi, Peter Stordy, and Antonio de la Vega de León. 2020. Integrating FATE/critical data studies into data science curricula: where are we going and how do we get there?. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 425–435. https://doi.org/10.1145/3351095.3372832

[7] Ruha Benjamin. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. John Wiley & Sons.

[8] Karl-Emi Kjær Bilstrup, Magnus Høholt Kaspersen, Ira Assent, Simon Enni, and Marianne Graves Petersen. 2022. From Demo to Design in Teaching Machine Learning. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 2168–2178. https://doi.org/10.1145/3531146.3534634

[9] Abeba Birhane. 2021. Algorithmic injustice: a relational ethics approach. *Patterns* 2, 2 (Feb. 2021), 100205. https://doi.org/10.1016/j.patter.2021.100205

[10] Frank S. Bloch (Ed.). 2011. *The global clinical movement: educating lawyers for social justice*. Oxford University Press, New York Oxford. https://doi.org/10.1093/acprof:oso/9780195381146.001.0001

[11] Chris Bopp, Ellie Harmon, and Amy Voida. 2017. Disempowered by Data: Nonprofits, Social Enterprises, and the Consequences of Data-Driven Work. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 3608–3619. https://doi.org/10.1145/3025453.3025694

[12] Robert R Bradley. 2015. *A quantitative study of the effects of service learning orientation classes on full-time, first-time freshmen*. Ed.D. Thesis. Tennessee State University. https://digitalscholarship.tnstate.edu/dissertations/AAI3723599

[13] Robert G. Bringle and Patti H. Clayton. 2012. Civic Education through Service Learning: What, How, and Why? In *Higher Education and Civic Engagement: Comparative Perspectives*, Lorraine McIlrath, Ann Lyons, and Ronaldo Munck (Eds.). Palgrave Macmillan US, New York, 101–124. https://doi.org/10.1057/9781137074829_7

[14] Joy Buolamwini and Timnit Gebru. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*. PMLR, 77–91. https://proceedings.mlr.press/v81/buolamwini18a.html

[15] Jenna Burrell and Marion Fourcade. 2021. The Society of Algorithms. *Annual Review of Sociology* 47, 1 (July 2021), 213–237. https://doi.org/10.1146/annurev-soc-090820-020800

[16] Francina Cantatore, Linda Crane, and Deborah Wilmoth. 2016. Defining Clinical Education: Parallels in Practice. *Australian Journal of Clinical Education* 1, 1 (Dec. 2016). https://doi.org/10.53300/001c.5087 Publisher: Bond University.

[17] Lauren M. Chambers. 2025. Beyond Big Tech: Advocacy Technologists within Mission-Driven Civil Society Organizations. *Proc. ACM Hum.-Comput. Interact.* 9, CSCW7 (Oct. 2025), 378:1–378:31. https://doi.org/10.1145/3757559

[18] Heewon Chang, Faith Ngunjiri, and Kathy-Ann C Hernandez. 2016. *Collaborative autoethnography*. Routledge.

[19] Erwin Chemerinsky. 2009. Why Not Clinical Education? *Clinical Law Review* 16, 1 (2009), 35–42. https://heinonline.org/HOL/P?h=hein.journals/clinic16&i=39

[20] Shreya Chowdhary, Sam Daitzman, Ruby Eisenbud, Emma Pan, and Erhardt Graeff. 2020. Care and Liberation in Creating a Student-Led Public Interest Technology Clinic. In *2020 IEEE International Symposium on Technology and Society (ISTAS)*. IEEE Press, Tempe, AZ, USA, 164–175. https://doi.org/10.1109/ISTAS50296.2020.9462188

[21] Randy W. Connolly. 2012. Is there service in computing service learning?. In *Proceedings of the 43rd ACM technical symposium on Computer Science Education (SIGCSE '12)*. Association for Computing Machinery, New York, NY, USA, 337–342. https://doi.org/10.1145/2157136.2157238

[22] Julian Thomas Costa. 2017. Service Learning, Project Management and Professional Development. *Business Education Innovation Journal* 9, 2 (Dec. 2017), 32–38.

[23] Sasha Costanza-Chock. 2020. *Design Justice: Community-led Practices to Build the Worlds we Need.* The MIT Press, Cambridge, Massachusetts.

[24] Sasha Costanza-Chock, Maya Wagoner, Berhan Taye, Caroline Rivas, Chris Schweidler, Georgia Bullen, and T4SJ Project. 2018. *#More Than Code: Practitioners reimagine the landscape of technology for justice and equity.* Technical Report. Research Action Design & Open Technology Institute. https://morethancode.cc/T4SJ_fullreport_082018_AY_web.pdf

[25] Roderic Crooks and Morgan Currie. 2021. Numbers will not save us: Agonistic data practices. *The Information Society* 37, 4 (Aug. 2021), 201–213. https://doi.org/10.1080/01972243.2021.1920081 Publisher: Routledge _eprint: https://doi.org/10.1080/01972243.2021.1920081.

[26] Shiva Darian, Aarjav Chauhan, Ricky Marton, Janet Ruppert, Kathleen Anderson, Ryan Clune, Madeline Cupchak, Max Gannett, Joel Holton, Elizabeth Kamas, Jason Kibozi-Yocka, Devin Mauro-Gallegos, Simon Naylor, Meghan O'Malley, Mehul Patel, Jack Sandberg, Troy Siegler, Ryan Tate, Abigil Temtim, Samantha Whaley, and Amy Voida. 2023. Enacting Data Feminism in Advocacy Data Work. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1 (April 2023), 47:1–47:28. https://doi.org/10.1145/3579480

[27] Clare Delany and Elizabeth Molloy. 2009. *Clinical Education in the Health Professions: An Educator's Guide.* Elsevier Health Sciences.

[28] Clare Delany and Elizabeth Molloy. 2009. Critical reflection in clinical education: beyond the "swampy lowlands". In *Clinical Education in the Health Professions: An Educator's Guide*, Clare Delany and Elizabeth Molloy (Eds.). Elsevier Health Sciences.

[29] Anke Dittmar. 2023. Moments of Surprise: Preparing Technology-Oriented Students for Interdisciplinary Work. In *Proceedings of the European Conference on Cognitive Ergonomics 2023 (ECCE '23)*. Association for Computing Machinery, New York, NY, USA, 1–8. https://doi.org/10.1145/3605655.3605673

[30] Jon Dubin. 1998. Clinical Design for Social Justice Imperatives. *SMU Law Review* 51, 5 (Jan. 1998), 1461. https://scholar.smu.edu/smulr/vol51/iss5/4

[31] J. Paul Dunne and Elisabeth Sköns. 2021. New technology and the U.S. military industrial complex. *The Economics of Peace and Security Journal* 16, 2 (Oct. 2021). https://doi.org/10.15355/epsj.16.2.5

[32] Carolyn Ellis, Tony E Adams, and Arthur P Bochner. 2011. Autoethnography: an overview. *Historical social research/Historische sozialforschung* (2011), 273–290.

[33] Severin Engelmann, Madiha Zahrah Choksi, Angelina Wang, and Casey Fiesler. 2024. Visions of a Discipline: Analyzing Introductory AI Courses on YouTube. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*. Association for Computing Machinery, New York, NY, USA, 2400–2420. https://doi.org/10.1145/3630106.3659045

[34] Virginia Eubanks. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor.* St. Martin's Publishing Group.

[35] Janet Eyler. 2002. Reflection: Linking Service and Learning—Linking Students and Communities. *Journal of Social Issues* 58, 3 (2002), 517–534. https://doi.org/10.1111/1540-4560.00274 _eprint: https://spssi.onlinelibrary.wiley.com/doi/pdf/10.1111/1540-4560.00274.

[36] Casey Fiesler, Natalie Garrett, and Nathan Beard. 2020. What Do We Teach When We Teach Tech Ethics? A Syllabi Analysis. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education (SIGCSE '20)*. Association for Computing Machinery, New York, NY, USA, 289–295. https://doi.org/10.1145/3328778.3366825

[37] Paulo Freire. 1968. *Pedagogy of the Oppressed.* Continuum.

[38] Shuang Frost, Aleksandrina Goeva, Javin Pombra, William Seaton, Sara Stoudt, Ana Trisovic, Chris Wang, and Catherine Zucker. 2020. Kaleidoscopic Perspectives on Practicum-Based Data Science Education. *Harvard Data Science Review* 3, 1 (Dec. 2020). https://doi.org/10.1162/99608f92.ba6336bc Publisher: The MIT Press.

[39] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hannah Wallach, Hal Duamé, and Kate Crawford. 2021. Datasheets for Datasets. *Commun. ACM* 64, 12 (Dec. 2021), 86–92. https://cacm.acm.org/magazines/2021/12/256932-datasheets-for-datasets/fulltext

[40] Erhardt Graeff and Alison Wood. 2021. Undergraduate Engineering as Civic Professionalism. *The Good Society* 30, 1-2 (Dec. 2021), 76–95. https://doi.org/10.5325/goodsociety.30.1-2.0076 Publisher: Duke University Press.

[41] Colin M. Gray and Shruthi Sai Chivukula. 2019. Ethical Mediation in UX Practice. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3290605.3300408

[42] Ben Green. 2021. Data Science as Political Action: Grounding Data Science in a Politics of Justice. *Journal of Social Computing* 2, 3 (Sept. 2021), 249–265. https://doi.org/10.23919/JSC.2021.0029

[43] bell hooks. 1994. *Teaching to transgress: education as the practice of freedom.* Routledge, Taylor & Francis Group, New York. https://doi.org/10.4324/9780203700280

[44] Lilly C. Irani and M. Six Silberman. 2016. Stories We Tell About Labor: Turkopticon and the Trouble with "Design". *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (May 2016), 4573–4586. https://doi.org/10.1145/2858036.2858592

[45] Abigail Z. Jacobs and Hanna Wallach. 2021. Measurement and Fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 375–385. https://doi.org/10.1145/3442188.3445901

[46] Peter M. Jansson, Ravi P. Ramachandran, John L. Schmalzel, and Shreekanth A. Mandayam. 2010. Creating an Agile ECE Learning Environment Through Engineering Clinics. *IEEE Transactions on Education* 53, 3 (Aug. 2010), 455–462. https://doi.org/10.1109/TE.2009.2027431 Conference Name: IEEE Transactions on Education.

[47] Yasmin Kafai, Jean Griffin, Quinn Burke, Michelle Slattery, Deborah Fields, Rita Powell, Michele Grab, Susan Davidson, and Joseph Sun. 2013. A cascading mentoring pedagogy in a CS service learning course to broaden participation and perceptions. In *Proceeding of the 44th ACM technical symposium on Computer science education (SIGCSE '13)*. Association for Computing Machinery, New York, NY, USA, 101–106. https://doi.org/10.1145/2445196.2445228

[48] Marcy L. Karin and Robin R. Runge. 2010. Toward Integrated Law Clinics That Train Social Change Advocates. *Clinical Law Review* 17, 2 (2010), 563–614. https://heinonline.org/HOL/P?h=hein.journals/clinic17&i=570

[49] Os Keyes, Jevan Hutson, and Meredith Durbin. 2019. A Mulching Proposal: Analysing and Improving an Algorithmic System for Turning the Elderly into High-Nutrient Slurry. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (CHI EA '19)*. Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3290607.3310433

[50] Mallory Knodel, Mallika Balakrishnan, and Lauren Chambers. 2025. Shaping a Profession, Building a Community: A Practitioner-Led Investigation of Public Interest Technologists in Civil Society.

[51] P. M. Krafft, Meg Young, Michael Katell, Jennifer E. Lee, Shankar Narayan, Micah Epstein, Dharma Dailey, Bernease Herman, Aaron Tam, Vivian Guetler, Corinne Bintz, Daniella Raz, Pa Ousman Jobe, Franziska Putz, Brian Robick, and Bissan Barghouti. 2021. An Action-Oriented AI Policy Toolkit for Technology Audits by Community Advocates and Activists. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 772–781. https://doi.org/10.1145/3442188.3445938

[52] Vito Laterza. 2021. Could Cambridge Analytica Have Delivered Donald Trump's 2016 Presidential Victory? An Anthropologist's Look at Big Data and Political Campaigning. *Public Anthropologist* 3, 1 (March 2021), 119–147. https://doi.org/10.1163/25891715-03010007

[53] Cindy Lin and Silvia Margot Lindtner. 2021. Techniques of Use: Confronting Value Systems of Productivity, Progress, and Usefulness in Computing and Design. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–16. https://doi.org/10.1145/3411764.3445237

[54] Michael Madaio, Shivani Kapania, Rida Qadri, Ding Wang, Andrew Zaldivar, Remi Denton, and Lauren Wilcox. 2024. Learning about Responsible AI On-The-Job: Learning Pathways, Orientations, and Aspirations. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*. Association for Computing Machinery, New York, NY, USA, 1544–1558. https://doi.org/10.1145/3630106.3658988

[55] Milagros Miceli, Adio-Adet Dinika, Krystal Kauffman, Camilla Salim Wagner, Laurenz Sachenbacher, Alex Hanna, and Timnit Gebru. 2025. Methodological Considerations for Centering Workers' Epistemic Authority in AI Research. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* 8, 2 (Oct. 2025), 1698–1710. https://doi.org/10.1609/aies.v8i2.36667

[56] Yeshimabeit Milner. [n. d.]. About D4BL. https://d4bl.org/about

[57] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)*. Association for Computing Machinery, New York, NY, USA, 220–229. https://doi.org/10.1145/3287560.3287596

[58] Jared Moore. 2020. Towards a more representative politics in the ethics of computer science. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 414–424. https://doi.org/10.1145/3351095.3372854

[59] Deirdre K. Mulligan, Joshua A. Kroll, Nitin Kohli, and Richmond Y. Wong. 2019. This Thing Called Fairness: Disciplinary Confusion Realizing a Value in Technology. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–36. https://doi.org/10.1145/3359221 arXiv: 1909.11869.

[60] Michael I. Norton, Daniel Mochon, and Dan Ariely. 2012. The IKEA effect: When labor leads to love. *Journal of Consumer Psychology* 22, 3 (July 2012), 453–460. https://doi.org/10.1016/j.jcps.2011.08.002

[61] Isak Nti Asare, Scott Shackleford, Jungwoo Chun, and Sarah Powazek. 2025. Protecting Communities while Training Future Cybersecurity Professionals: Lessons from the Consortium of Cybersecurity Clinics. *The Cyber Defense Review* 10, 2 (Dec. 2025), 241–265. https://doi.org/10.55682/cdr/1dm8-wekj

[62] Ziad Obermeyer and Sendhil Mullainathan. 2019. Dissecting Racial Bias in an Algorithm that Guides Health Decisions for 70 Million People. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)*. Association for Computing Machinery, New York, NY, USA, 89. https://doi.org/10.1145/3287560.3287593

[63] Jamie Payton, Tiffany Barnes, Kim Buch, Audrey Rorrer, and Huifang Zuo. 2015. The effects of integrating service learning into computer science: an inter-institutional longitudinal study. *Computer Science Education* 25, 3 (July 2015), 311–324. https://doi.org/10.1080/08993408.2015.1086536 Publisher: Routledge _eprint: https://doi.org/10.1080/08993408.2015.1086536.

[64] Lucy Pei and Roderic Crooks. 2023. Grassroots Data Activism. *JCMS: Journal of Cinema and Media Studies* 62, 4 (2023), 188–192. https://muse.jhu.edu/pub/349/article/904636 Publisher: Michigan Publishing.

[65] Lucy Pei, Benedict Salazar Olgado, and Roderic Crooks. 2022. Narrativity, Audience, Legitimacy: Data Practices of Community Organizers. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. ACM, New Orleans LA USA, 1–6. https://doi.org/10.1145/3491101.3519673

[66] Jeria Quesenberry, Randy Weinberg, and Larry Heimann. 2012. Experiences in service-learning pedagogy: lessons for recruitment and retention of under represented groups. In *Proceedings of the 50th annual conference on Computers and People Research (SIGMIS-CPR '12)*. Association for Computing Machinery, New York, NY, USA, 89–90. https://doi.org/10.1145/2214091.2214116

[67] Fran Quigley. 1995. Seizing the Disorienting Moment: Adult Learning Theory and the Teaching of Social Justice in Law School Clinics. *Clinical Law Review* 2, 1 (1995), 37–72.

[68] Inioluwa Deborah Raji, Morgan Klaus Scheuerman, and Razvan Amironesei. 2021. You Can't Sit With Us: Exclusionary Pedagogy in AI Ethics Education. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 515–525. https://doi.org/10.1145/3442188.3445914

[69] Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 33–44. https://doi.org/10.1145/3351095.3372873

[70] Sarah T. Roberts. 2021. *Behind the Screen: Content Moderation in the Shadows of Social Media* (reprint edition ed.). Yale University Press.

[71] Fujiko Robledo Yamamoto, Lecia Barker, and Amy Voida. 2023. CISing Up Service Learning: A Systematic Review of Service Learning Experiences in Computer and Information Science. *ACM Trans. Comput. Educ.* 23, 3 (Sept. 2023), 37:1–37:56. https://doi.org/10.1145/3610776

[72] Fujiko Robledo Yamamoto, Lecia Barker, and Amy Voida. 2024. Faculty, Student, and Community Partner Experiences in Computer and Information Science Service Learning. *ACM Trans. Comput. Educ.* 24, 3 (May 2024), 31:1–31:26. https://doi.org/10.1145/3654678

[73] Adrian Rusu and Amalia Rusu. 2023. Fostering the Innovative Mindset: Entrepreneurship Clinic Model for Computer Science Students. In *Proceedings of the 2023 Conference on Innovation and Technology in Computer Science Education V. 1 (ITiCSE 2023)*. Association for Computing Machinery, New York, NY, USA, 464–470. https://doi.org/10.1145/3587102.3588812

[74] Dan Schiller. 1994. From culture to information and back again: Commoditization as a route to knowledge. *Critical Studies in Mass Communication* 11, 1 (March 1994), 92–115. https://doi.org/10.1080/15295039409366887 Publisher: Routledge _eprint: https://doi.org/10.1080/15295039409366887.

[75] Aure Schrock. 2024. *Politics Recoded: The Infrastructural Organizing of Code for America*. MIT Press.

[76] Andrew D. Selbst, Danah Boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)*. Association for Computing Machinery, New York, NY, USA, 59–68. https://doi.org/10.1145/3287560.3287598

[77] Hong Shen, Wesley H. Deng, Aditi Chattopadhyay, Zhiwei Steven Wu, Xu Wang, and Haiyi Zhu. 2021. Value Cards: An Educational Toolkit for Teaching Social Impacts of Machine Learning through Deliberation. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 850–861. https://doi.org/10.1145/3442188.3445971

[78] Katie Shilton. 2013. Values Levers: Building Ethics into Design. *Science, Technology, & Human Values* 38, 3 (May 2013), 374–397. https://doi.org/10.1177/0162243912436985

[79] Katie Shilton. 2018. Engaging Values Despite Neutrality: Challenges and Approaches to Values Reflection during the Design of Internet Infrastructure. *Science, Technology, & Human Values* 43, 2 (March 2018), 247–269. https://doi.org/10.1177/0162243917714869

[80] Herbert A. Simon. 1955. A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics* 69, 1 (Feb. 1955), 99. https://doi.org/10.2307/1884852

[81] Herbert A. Simon. 1991. Bounded Rationality and Organizational Learning. *Organization Science* 2, 1 (Feb. 1991), 125–134. https://doi.org/10.1287/orsc.2.1.125

[82] Herbert A. Simon. 1996. *The sciences of the artificial* (3rd ed ed.). MIT Press, Cambridge, Mass. original-date:1969.

[83] Lucy Suchman, Sireesh Gururaja, and David Gray Widder. 2025. AI interdisciplinarity as critical technical practice. *Cambridge Forum on AI: Culture and Society* 1 (Jan. 2025), e2. https://doi.org/10.1017/cfc.2025.10004

[84] Cella M. Sum, Jiayin Zhi, Amil N.T. Cook, Patrick James Cooper, Arturo Lozano, Tj Johnson, Jason Perez, Rayid Ghani, Michael Skirpan, Motahhare Eslami, Hong Shen, and Sarah E. Fox. 2025. "You're in a Ferrari. I'm Waiting for the Bus": Confronting Tensions in Community-University Partnerships. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (May 2025), 1–28. https://doi.org/10.1145/3711110

[85] Harini Suresh and John Guttag. 2021. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. In *Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '21)*. Association for Computing Machinery, New York, NY, USA, 1–9. https://doi.org/10.1145/3465416.3483305

[86] Stephanie Tena-Meza, Miroslav Suzara, and Aj Alvero. 2022. Coding with Purpose: Learning AI in Rural California. *ACM Trans. Comput. Educ.* 22, 3 (Sept. 2022), 24:1–24:18. https://doi.org/10.1145/3513137

[87] Matthew Ung, Andrew Lin, and Daniel Votipka. 2024. WiP: A Qualitative Study of Service-Learning Oriented Cybersecurity Clinics' Processes and Challenges. *USENIX Symposium on Usable Privacy and Security (SOUPS)* (Aug. 2024).

[88] Hanna Wallach, Meera Desai, A. Feder Cooper, Angelina Wang, Chad Atalla, Solon Barocas, Su Lin Blodgett, Alexandra Chouldechova, Emily Corvi, P. Alex Dow, Jean Garcia-Gathright, Alexandra Olteanu, Nicholas J. Pangakis, Stefanie Reed, Emily Sheng, Dan Vann, Jennifer Wortman Vaughan, Matthew Vogel, Hannah Washington, and Abigail Z. Jacobs. 2025. Position: Evaluating Generative AI Systems Is a Social Science Measurement Challenge. In *Proceedings of the 42nd International Conference on Machine Learning*. PMLR, 82232–82251. https://proceedings.mlr.press/v267/wallach25a.html ISSN: 2640-3498.

[89] Anne L. Washington and Joanne Cheung. 2024. Public Interest. In *Keywords of the Datafied State*, Jenna Burrell, Ranjit Singh, and Patrick Davison (Eds.). Data & Society, 93–105. https://www.ssrn.com/abstract=4734250

[90] Alicia Nicki Washington. 2020. When Twice as Good Isn't Enough: The Case for Cultural Competence in Computing. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education (SIGCSE '20)*. Association for Computing Machinery, New York, NY, USA, 213–219. https://doi.org/10.1145/3328778.3366792

[91] Tom Wheeler. 2023. *Techlash: Who Makes the Rules in the Digital Gilded Age?* (1 ed.). Rowman & Littlefield Publishers, Blue Ridge Summit.

[92] Cedric Whitney. 2025. Contextual Evaluations. In *From Data to Evaluation: Investigating the Limits of the Individual-Control Model in AI Governance*. UC Berkeley. https://escholarship.org/uc/item/7135b09g ISBN: 9798288862052.

[93] David Gray Widder, Derrick Zhen, Laura Dabbish, and James Herbsleb. 2023. It's about power: What ethical concerns do software engineers have, and what do they (feel they can) do about them?. In *2023 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Chicago IL USA, 467–479. https://doi.org/10.1145/3593013.3594012

[94] Langdon Winner. 1980. Do Artifacts Have Politics? *Daedalus* 109, 1 (1980). https://www.jstor.org/stable/20024652

[95] Stephen Wizner. 2000. Beyond Skills Training Papers Presented at the Rutgers-Newark Law School Conference on the Social Justice Mission of Clinical Education. *Clinical Law Review* 7, 2 (2000), 327–340. https://heinonline.org/HOL/P?h=hein.journals/clinic7&i=333

[96] Richmond Y. Wong. 2021. Tactics of Soft Resistance in User Experience Professionals' Values Work. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 1–28. https://doi.org/10.1145/3479499

[97] Richmond Y. Wong and Tonya Nguyen. 2021. Timelines: A World-Building Activity for Values Advocacy. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3411764.3445447

[98] Carter Godwin Woodson. 2009. *The Mis-Education of the Negro*. The Journal of Pan African Studies. https://jpanafrican.org/ebooks/3.4eBookThe%20Mis-Education.pdf original-date:1933.

[99] Yaxuan Yin, Shamya Karumbaiah, and Shona Acquaye. 2025. Responsible AI in Education: Understanding Teachers' Priorities and Contextual Challenges. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*. Association for Computing Machinery, New York, NY, USA, 2705–2727. https://doi.org/10.1145/3715275.3732176

[100] Patrick L. Yorio and Feifei Ye. 2012. A meta-analysis on the effects of service-learning on the social, personal, and cognitive outcomes of learning. *Academy of Management Learning & Education* 11, 1 (2012), 9–27. https://doi.org/10.5465/amle.2010.0072 Place: US Publisher: Academy of Management.

[101] Shoshana Zuboff. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (1st edition ed.). PublicAffairs, New York.

Received 3 January 2026; accepted 2 March 2026