

A Critical Pragmatism Approach for AI Ethics: Lessons from Urban Planning Theory

Jennah Gosciak^{1,2}, Karen Levy², Allison Koenecke^{1,2}

¹Cornell Tech

²Cornell University

jrg377@cornell.edu, karen.levy@cornell.edu, koenecke@cornell.edu

Abstract

As data scientists grapple with increasingly complex ethical decisions in machine learning (ML) and algorithm design, scholarship on AI ethics has offered multiple solutions, from formal mathematical definitions of fairness to holistic fairness frameworks drawn from various academic disciplines. Still, developing ethical and fair algorithmic systems in practice remains an ongoing challenge. In this paper, we draw a parallel between problems in algorithmic fairness and urban planning. We frame algorithmic fairness problems as ‘wicked problems,’ a term originating from the planning and policy space to describe the intractable, value-laden, and complex nature of this work. As such, we argue that the field of AI ethics can learn from theoretical work in urban planning in ameliorating its own set of wicked problems. Urban planning is typically concerned with practical issues of governance, resource allocation, stakeholder engagement, and conflicts involving deep-seated differences. These are challenges that purely technical approaches in AI easily overlook. We present a flexible set of recommendations based on the urban planning theory approach of *critical pragmatism* — a reflective and deliberative approach to addressing wicked problems that considers what practitioners actually *do* in the face of conflict and power. We apply our recommendations to several case studies: automated mortgage lending, school choice, and femicide counterdata collection. Researchers and practitioners can incorporate these recommendations derived from urban planning into their ongoing work to more holistically address problems related to AI ethics.

1 Introduction

Understanding the social and ethical impact of algorithms is complex and has been the subject of substantial interdisciplinary research with contributions from computer science, law, philosophy, and the social sciences (Chouldechova 2017; Kleinberg, Mullainathan, and Raghavan 2016; Jacobs and Wallach 2021; Binns 2018; Heidari et al. 2019; Arif Khan, Manis, and Stoyanovich 2022; Selbst et al. 2019; Davis, Williams, and Yang 2021). In particular, there is a growing body of theoretical work that takes a critical and holistic view of algorithmic fairness – from data feminism (D’Ignazio and Klein 2023) and non-ideal algorithmic fairness (Fazelpour and Lipton 2020) to algorithmic re-

alism (Green and Viljoen 2020) and pipeline-aware fairness (Black et al. 2023). Implementing the guidance from these holistic approaches in real-world settings, however, remains a challenge. Recent work has highlighted the persistent barriers facing industry practitioners and has called for a more pragmatic and practical lens (Holstein et al. 2019; Lee and Singh 2021; Deng et al. 2022; Smith et al. 2025; Deng, Barocas, and Wortman Vaughan 2025).

Our work fills this gap by providing a series of practical recommendations that data scientists and engineers can use. We develop these recommendations by comparing *algorithmic fairness problems* to similar problems in *urban planning*.¹ Algorithmic fairness problems relate to the complex ethical trade-offs and governance concerns involved in designing algorithmic systems. Planning, broadly, refers to a collection of activities associated with organizing social and political life. Planning problems can be local or global; for example, economic, environmental, and regional planning are also types of planning fields (Friedmann 1987).

To concretize our analogy, we use the framing of “wicked problems” — a term proposed by planning and design theorists Horst Rittel and Melvin Webber (1973).² Wicked problems are complicated, value-laden, and subject to reasonable disagreement. The widely influential term has been cited in fields like operations research, software engineering, and human-computer interaction. However, fundamentally, it is a term from the planning literature.

Framing algorithmic fairness problems as “wicked problems” allows us to engage in translational work, borrowing ideas from urban planning theorists who have sought to resolve similarly complex “wicked” problems for decades. Much of the recent theoretical work addressing wicked problems in urban planning has focused on *practical experience* and *philosophical pragmatism*. Planning theorists seek to

¹We use the term “urban planning,” but this literature is also related to policy analysis and public administration (Forester 1981).

²Though the term wicked problems has been widely cited across numerous fields, both Rittel and Webber were originally speaking to a planning audience. The two worked on research related to design thinking, architecture, and city and regional planning. Horst Rittel was a professor in the Department of Architecture at U.C. Berkeley (Churchman et al. 2007); Melvin Webber was faculty in the Department of City and Regional Planning at U.C. Berkeley (Maclay 2006).

guide practitioners in what they can actually *do* to resolve conflict and navigate value differences. We argue that there are similar shifts in algorithmic fairness, analogous to the evolution of pragmatic thinking in the planning literature. To address algorithmic fairness problems *in practice*, we believe that scholars can learn from the ways in which urban planners have navigated the messy and political reality of their work for many decades.

Critical pragmatism attempts to resolve the practice-theory tensions that Rittel and Webber (1973) first identified with “wicked problems.” It draws on prior theoretical work in urban planning and is grounded in empirical insights (e.g., first-hand accounts from urban planners). It addresses planning theory’s “dirty little secret” – that even, despite extensive theorizing, “somebody still has to do the work” (Forester 2017).

Our primary contribution in this paper is to propose a set of four recommendations based on the urban planning theory approach of critical pragmatism. We recommend that practitioners: (1) reflect-in-action, (2) hold public meetings that involve deliberation, (3) listen critically and seek out creative negotiations, and (4) share practice-oriented stories. We apply our four recommendations to three case studies: mortgage lending, school choice, and activism against femicide. In each case study, we describe the current algorithm and context. We then discuss how our recommendations might improve on the current approach and where they fall short.

Our recommendations complement existing scholarship in algorithmic fairness and AI ethics, which has similarly called for increases in participation and critiqued power (Suresh et al. 2022; Delgado et al. 2023; Sloane et al. 2022; Kulynych et al. 2020; Suresh et al. 2024; Costanza-Chock 2020; D’Ignazio and Klein 2023). Still, our recommendations formalize and advance important points that we believe are under-emphasized in the literature.

First, several of our recommendations are aimed at individuals, and can be utilized without institutional resources. These recommendations are useful for generalizing across the diverse and context-specific problems that practitioners often encounter. Second, critical pragmatism emphasizes important social and dialogical tools needed to adequately address AI ethics concerns. Participation is useful not as a co-design tool to elicit information, but to advance a democratic and deliberative process, mediating across seemingly intractable value conflicts and building trust. This perspective builds on philosophers like Jürgen Habermas and Paulo Freire (Forester 2017), whose work on communication and education is less commonly cited in AI ethics literature.

To underscore this point, we distinguish *deliberation* from *participation*. Deliberation requires active engagement from stakeholders (similar to the kinds of listening sessions, working groups, and public meetings that are common in urban planning, but less so in conventional approaches to algorithm design or ML). It is not that these technical processes never involve participation from stakeholders, but in many settings, there is little emphasis on communication – *listening* in particular – as a tool for resolving conflict and trade-offs. As Delgado et al. (2023) note, many

participatory ML projects consult with stakeholders only once, often through surveys or interviews. Sustained engagement involving genuine deliberation is relatively rare. Still, in several instances where researchers have reflected on effective practices (e.g., in ML or AI for social good initiatives), such reflections have revealed the following elements of success: the sensitive manner in which information was communicated, the gradual building of trust and accountability through carefully structuring the participatory process (Sendak et al. 2020), and “real listening” (Shi et al. 2022). **We argue that these elements should be more clearly articulated as guidance for practitioners aiming to develop fair and ethical algorithmic tools.**

This paper is intentionally theoretical in scope. We do not claim that critical pragmatism will *solve* all problems in AI ethics. In fact, the premise of our argument rests on the point that such problems are not solvable. Instead, by engaging with the urban planning literature, we offer tools that practitioners can use to incrementally forge a path forward.

The rest of the paper is structured as follows. In Section 2, we define urban planning and introduce the concept of wicked problems. We describe the evolution of pragmatic thinking in urban planning in Section 3. Then, in Section 4, we introduce the framework of critical pragmatism, discuss foundational concepts, and note relevant critiques. In Section 5, we turn to recent work in algorithmic fairness and justify the analogy between the two fields. We briefly summarize different algorithmic fairness approaches, categorizing them as either idealized or holistic. We then discuss ongoing practical concerns with algorithmic fairness related to AI ethics and governance. In Section 6, drawing on critical pragmatism, we provide four overarching recommendations that scholars and practitioners can flexibly use in their ongoing work. In Section 7, we apply these recommendations to several case studies. We conclude by encouraging scholars and practitioners in AI ethics to consider these recommendations drawn from urban planning theory.

2 Defining Wicked Problems

Urban planning stems from the broader, interdisciplinary field of planning. Planning has been called the link between knowledge and action (Friedmann 1987), the “guidance of future action” (Forester 1999), or the collection of a diverse set of practices (Alexander 2022). For many decades, scholars have debated different definitions (Alexander 1981). Planners navigate high-stakes political conflicts and guide the design, implementation, and impact of policies with evidence-backed decision-making and stakeholder engagement. The field is also closely related to architecture and design. Planning is: “societal, future-orientated, non-routinised, deliberate, strategic and linked to action” (Alexander 1981).

Planning also deals with “wicked problems” —intractable and ill-formed problems that stand in contrast to the “tame” or “benign” problems common in science and engineering (Rittel and Webber 1973) (see Figure 1). Examples of wicked problems include serious policy issues with no straightforward solutions like poverty or climate change.




Figure 1: Connections are depicted between urban planning theory and AI ethics. Rational-comprehensive planning and idealized fairness definitions are more apt for “tame” or “benign” problems, as discussed in Rittel and Webber (1973), while pragmatic planning and holistic fairness address “wicked problems.” These terms are defined in Sections 2, 3, and 5.2.

To provide a concrete example, consider housing affordability – a concern in many cities around the world. Affordable housing is a central issue in urban planning, and one of the most common planning specializations according to the American Planning Association (2025). In the United States (U.S.), there is currently no state or county where a renter who works full-time at minimum wage can afford a two-bedroom apartment (National Low Income Housing Coalition 2023). Yet we know that housing access and stability affect life outcomes like employment, physical and mental well-being, educational achievement, and financial success (Desmond and Gershenson 2016; Gubits et al. 2015).

Housing affordability meets the 10-step criteria that Rittel and Webber (1973) use to define wicked problems. For illustrative purposes, we discuss the first one: “there is no definitive formulation” of the problem. The other steps are discussed in more detail in the appendix.³ This first step is true because there are different explanations for why housing is unaffordable. Some organizers and advocates call for more public housing (Fischer, Acosta, and Bailey 2021). Others argue for housing vouchers (Bailey 2022). To some, zoning is the issue: upzoning neighborhoods should produce denser housing and reduce rents. Calls for upzoning are often met with resistance, due to fears about neighborhood change or displacement (Davis 2021). In land use decisions, affordable housing is often the site of intense conflict, with homeowners, renters, housing advocates, and government officials on very different sides of the issue.

3 Urban Planning History Addressing Wicked Problems

Rittel and Webber (1973)’s “wicked problems” was a reaction to the emphasis on scientific and rational thinking in modern professions around the mid-20th century. This thinking produced the *rational-comprehensive* model of planning (Whittemore 2015) in which planners functioned much like scientific experts. Rational planning has been widely criticized with many real examples of failure in U.S. history (Lindblom 1959; Schön 2016; Scott 1998).

Urban renewal stands out as one of the most emblematic and problematic examples. Shortly after World War II, the U.S. federal government sought to revitalize cities through redevelopment. At the time, rational-comprehensive plan-

ning dominated. Systems analysis, computational models, and quantitative analysis were the primary methods of urban planners (Hall 1989). Urban renewal ushered in a period of widespread slum clearance and new construction that left a legacy of destruction (Gans 1959; Anderson 1964; Jacobs 1992). As a result, local governments eliminated entire neighborhoods and displaced low-income residents of color (Hall 1989; Von Hoffman 2008; Gans 1959).

Urbanists frequently cite Jane Jacobs as one of the central figures opposing urban renewal (Scott 1998). Jacobs saw the city as an “immense laboratory of trial and error” (Jacobs 1992). She criticized mid-20th century planners like Le Corbusier whose “City in the Park” style of architecture became the blueprint for many U.S. housing projects (Hunt 2001) and Robert Moses who executed large-scale infrastructure projects in New York City (NYC), often at the expense of the residents living there. These planners had abstract, utopian visions of cities, divorced from social and historical context (Callahan and Ikeda 2014; Jacobs 1992). Jacobs condemned numerous examples of such planning failures: the decline of the Morningside Heights neighborhood in NYC following slum clearance and the development of a middle-income housing project, underutilized and run-down civic centers built according to the ideals of the modernist City Beautiful movement like the Benjamin Franklin Parkway in Philadelphia, or the homogeneity of single-use areas like Wilshire Boulevard in Los Angeles (Jacobs 1992).

3.1 The Ensuing Turn to Pragmatism

The failures of rational planning are a warning to other fields. Rational planning had real and disastrous consequences on cities. However, these failures also ushered in a shift toward pragmatism in urban planning (see Figure 2). Rittel and Webber’s concept of “wicked problems” aligns with the onset of this shift. While influential across many fields, in urban planning their work has inspired newer generations of scholars to advance theory grounded in philosophical pragmatism (Forester 2020).

Pragmatism is a philosophical tradition that argues against pre-determined *a priori* beliefs and principles (Healey 2008). Pragmatists prioritize practical experience. They believe that hypotheses about what works in the world should be continuously discovered and tested empirically (Healey 2008). The rise of pragmatic thinking can be traced to three central figures: William James, Charles Sanders Peirce, and

³Manuscript with appendices: arxiv.org/abs/2605.02773.

John Dewey (Bacon 2012). Dewey in particular had a lasting influence on U.S. policy. His emphasis on experimentation, inquiry, and social justice influenced ideas on participatory democracy and popular support for social programs like the New Deal (Healey 2008). Dewey (1927) writes:

When we say that thinking and beliefs should be experimental, not absolutistic, we have then in mind a certain logic of method, not, primarily, the carrying on of experimentation like that of laboratories... Differences of opinion in the sense of differences of judgment as to the course which it is best to follow, the policy which it is best to try out, will still exist. But opinion in the sense of beliefs formed and held in the absence of evidence will be reduced in quantity and importance. No longer will views generated in view of special situations be frozen into absolute standards and masquerade as eternal truths.

Planning theorists have increasingly embraced pragmatic thinking—both older ideas from James, Peirce, and Dewey and more contemporary work from philosophers like Jürgen Habermas, Richard Bernstein, and Richard Rorty (Healey 2008). Examples of pragmatism in planning include: Rittel and Webber (1973)’s critiques about viewing policy and planning professions as a science, Lindblom (1959)’s work on incrementalism and experimentation, Friedmann (2008)’s transactive planning approach to generate dialogue and navigate difference, Davidoff (1965)’s notion of plural plans and advocacy planning model, and Forester (2013)’s *critical pragmatism*. We provide definitions for these early pragmatic approaches in the appendix. We focus on later approaches (reflective practice, deliberative practice, and critical pragmatism) in the following section.

4 Defining Critical Pragmatism

Critical pragmatism is a unifying framework that builds on pragmatic planning approaches. Proposed by theorist John Forester, critical pragmatism is concerned with practice, not ideals. It focuses on real-world examples through interviews and case studies. It also grapples with deliberation and disagreement in public settings (Forester 2013).

There are five key characteristics. First, it considers both process and outcomes, in part because of the “co-constructed” nature of how it approaches planning practice. Second, through its appreciation for different forms of knowledge, it encourages listening and interpretation. It considers the influence of power on relationships and claims about knowledge. Third, it is concerned with modes of communication and deliberation. Fourth, it offers lessons about process design and creative problem-solving. Fifth, it can help us shift “from a deconstructive skepticism toward a reconstructive imagination” (Forester 2013). Critical pragmatism does not force us to pessimistically accept reality as it is (as a simple “pragmatic” approach might), but remains optimistic that there are creative, alternative paths forward.

Many insights in critical pragmatism arise directly from analyzing stories about planners’ actual experiences. Forester believes that planners learn through storytelling with friends and colleagues – “practice stories” (Forester

1999). Drawing on the work of contemporary pragmatists like Habermas (1979) and Freire’s emphasis on dialogue and critical pedagogy (1970), Forester connects pragmatic methods to learning in groups and democratic deliberation. He argues for an active and interpretive stance to communication—what he calls “hermeneutic praxis” (Forester 2013, 2017). This method involves understanding how meaning and information shape action (Forester 1982) — similar to an interpretive and dialogical approach to policy analysis (Wagenaar 2014). Forester (2013) writes:

The critical pragmatist...will attend not only to the consequences of framing a problem in a certain way, but to the contingencies of the relations of power and authority that can make alternative frames and knowledge claims more or less plausible in the first place.

4.1 Reflective and Deliberative Practice

Two concepts are foundational to critical pragmatism: reflective and deliberative practice. In this section, we define both separately. However, we will broadly refer to them together as the two concepts go hand in hand.

Reflective Practice Forester’s critical pragmatism builds on Donald Schön (2016)’s book *The Reflective Practitioner*, which was first published in 1983. Schön’s concept of the “reflective practitioner” influenced many professional and design-oriented fields. This theory takes inspiration from earlier work in pragmatism, like Dewey’s ideas about inquiry and trial-and-error. Schön’s legacy was to contribute a new approach to the “epistemology of practice”—rooted in what practitioners actually do and dissolving the sharp boundary between theory and practice. Schön (2016) writes:

When someone reflects-in-action, he becomes a researcher in the practice context. He is not dependent on the categories of established theory and technique, but constructs a new theory of the unique case... Because his experimenting is a kind of action, implementation is built into his inquiry. Thus reflection-in-action can proceed, even in situations of uncertainty or uniqueness...”

Schön (2016) describes a series of steps to effectively learn and generalize from the many context-specific situations that arise in practice: (1) try one approach to framing the problem; (2) test this framing and consider any challenges that arise; (3) revise the initial framing of the problem by reflecting on the results—a process known as “back-talk” (i.e., unexpected results). While each problem may be unique, through this approach, one can eventually develop a “repertoire” of similar experiences to inform problems in new situations (Schön 2016).

Deliberative Practice Forester finds Schön’s image of the individual practitioner engaged in reflective practice too limiting (Schön 2016; Forester 2013). He expands on the notion of reflective practice by considering social and political context, multi-stakeholder settings, democratic and deliberative processes, and learning in groups. Thus, we can view Forester’s deliberative practitioner as an evolution of Schön’s reflective practitioner.

Deliberation is distinct from participation. Forester uses deliberation in the political sense, to denote group learning and dialogue. This definition aligns with contemporary pragmatists who connect the pragmatic emphasis on inquiry and problem-solving with collective deliberation (Anderson 2020; Misak 2004).

Both Schön (2016)'s notion of "reflection-in-action" and Forester (2013)'s deliberative practitioner offer ways to tackle wicked problems. Critical pragmatism additionally incorporates ideas about how practitioners can navigate communication amidst conflict, value differences, and power. We now discuss some critiques of critical pragmatism before turning back to a discussion of theoretical work in algorithmic fairness and comparing the two fields.

4.2 Critiques and Limitations

While Forester's work on critical pragmatism is included in many urban planning curricula, it has also received criticism, often for its overarching focus on communication. Communicative planning, exemplified by work from Patsy Healey (1992) and Judith Innes (1995), emphasizes consensus-building, collaboration, and dialogue. Critics of communicative planning raise several issues: (1) indifference to structural power dynamics (Forester's work is perhaps a counterexample (Forester 1980, 1989), though some might argue that he does not thoroughly consider structural dimensions of power); (2) the ambiguous definition of "planners" that assumes neutrality and benevolence (Allmendinger and Tewdwr-Jones 2002); (3) the assumption that consensus is desirable (Huxley 2000); and (4) the dominance of Western ideals and theories, which casts planning as both universal and unproblematic (Huxley and Yiftachel 2000). Many of these critiques overlap with broader concerns about the role of planning in upholding state power, and relate to criticism of more foundational theorists, such as Habermas (Huxley and Yiftachel 2000).

We encourage readers to consider these critiques and use them to frame similar discussions in AI ethics. Critical pragmatism advances a set of methods that can provide practical, incremental (Lindblom 1959) guidance about how to move forward. However, we do not suggest that critical pragmatism is a catch-all solution to any AI ethics problem, or even broader societal issues. Alternative approaches (e.g., grounded in political economy and critical theory) may be more useful in some settings (Huxley and Yiftachel 2000; Yiftachel and Huxley 2000).

5 Connections between Urban Planning Theory and AI Ethics

As algorithmic systems influence many high-stakes, societal decisions, the framing of wicked problems has become increasingly relevant to the field of AI ethics. Algorithms often must confront the underlying social and political dynamics that make such decisions inherently difficult, even in non-AI settings. Recent examples include: predictive policing (Lum and Isaac 2016; Green 2019), AI hiring (Fabris et al. 2025; Robinson 2022; Gerchick et al. 2025), and school choice algorithms (Colin Lecher 2021; Maddy Varner

2021). Competing definitions of fairness and equity further complicate these decisions. To illustrate the connection to "wicked problems," we provide a worked example in the appendix related to mortgage lending algorithms.

The overlap between urban planning and AI ethics – both grapple with "wicked problems" – suggests there may be value in learning from urban planning theory. Though prior work in AI ethics has cited Rittel and Webber (1973),⁴ the discussion of "wicked problems" in the literature has been relatively brief and the origins of the term within urban planning remain unexplored. The connection between the two fields is especially relevant because, as we discuss in the next section, the evolution of theoretical scholarship in AI ethics over the last decade mirrors that of urban planning.

5.1 The Comparison between Fields

We recognize that not all algorithmic fairness problems may compare to urban planning, but there are several reasons why the comparison between urban planning and algorithmic fairness is justified and even useful. First, as a scoping decision, we primarily focus on allocative harms in algorithmic fairness, as opposed to representational harms (Barocas, Hardt, and Narayanan 2023). Allocation decisions are widespread in both fields.

Second, urban planning is often discussed in the context of local (city or neighborhood-level) problems. This perspective overlooks the fact that urban planning involves local responses to broader societal issues, much like how data scientists and engineers may need to adapt scalable technologies to specific contexts. While *urban* planning deals with small-scale problems, the field of planning encompasses regional, economic, and environmental planning, which can all occur at non-local scales and involve extensive coordination (e.g., a national program). Algorithmic fairness problems similarly can vary in scale. We believe our recommendations extend even to settings involving large-scale technology. Participatory processes can reasonably occur even at the national level (e.g., see *Voices in the Code* (Robinson 2022)). Decisions about which stakeholders to include, communication choices, and problem framing may be even more critical and high stakes in such settings.

Third, many view urban planning as a field involving public-sector stakeholders, who may already approach "wicked problems" with a normative frame of justice and equity. While many urban planners do work in the public sector, urban planning encompasses a diverse set of roles in both public (e.g., local government) and private organizations (e.g., development companies), similar to roles in data science and engineering. We believe our recommendations, especially those aimed at individuals, are relevant, irrespective of a specific industry or role.

Lastly, we acknowledge that our recommendations may not be appropriate in some AI contexts where priorities are more narrowly tied to specific, private business interests. Our recommendations may be most effective when there is at least some commitment to developing fairer algorithms

⁴See prior work like Dobbe, Krendl Gilbert, and Mintz (2021); Scantamburlo (2021); Strauß (2021); Green (2019).

and resolving conflict, but disagreement about how. Additionally, even privately developed AI tools can have broader societal impact that warrants regulation. NYC Local Law 144 (The New York City Council 2021) is one example.

5.2 Algorithmic Fairness Approaches Mirror Urban Planning History

Borrowing terminology from Barocas, Hardt, and Narayanan (2023); Fazelpour and Lipton (2020), we classify two distinct approaches to fairness that have emerged in the literature on AI ethics: idealized vs. holistic fairness. This classification captures ongoing tensions in the field (see Figure 1): the limitations of idealized abstractions in comparison to the broader contextual and philosophical understandings of fairness (Green 2022). **The difficulty scholars face in reconciling different fairness definitions parallels the same dilemma Rittel and Webber confronted when proposing wicked problems.**

Idealized Fairness Idealized fairness definitions often formalize different fairness criteria mathematically and make simplifying assumptions that may be unrealistic. These approaches are precise and tractable, but may be too narrow. In real-world settings, they can lead to unfair outcomes (Green and Viljoen 2020; Green 2020; Kalluri 2020; Chohlas-Wood et al. 2023). For example, as researchers have shown, satisfying fairness via multiple formal definitions in many real-world instances can even be *impossible* both theoretically and in practice (Chouldechova 2017; Kleinberg, Mullainathan, and Raghavan 2016). Many established fairness frameworks and toolkits tend to use these methods (e.g., the *fairlearn* library (Weerts et al. 2023)).

Holistic Fairness In reaction to the limitations of idealized fairness approaches, researchers have turned to more holistic frameworks, drawing on scholarship from science and technology studies (STS), law, philosophy, and economics (Pessach and Shmueli 2023; Jacobs and Wallach 2021; Binns 2018; Heidari et al. 2019; Arif Khan, Manis, and Stoyanovich 2022; Lee, Floridi, and Singh 2021; Selbst et al. 2019; Davis, Williams, and Yang 2021). We classify approaches as holistic if they equate fairness with more expansive notions of equality and justice.

Guidance on implementing holistic fairness in practice, however, is not clear. Many holistic approaches treat fairness as context-dependent or as a social and political construct. These methods still require value judgments and lack a principled process. For example, operationalizing equal opportunity doctrines from political philosophy *in practice* is still far from straightforward (Arif Khan, Manis, and Stoyanovich 2022; Heidari et al. 2019). Likewise, Selbst et al. (2019)’s work drawn from STS theory argues for prioritization of social context and process over outcomes, which requires thoughtful and subjective judgment.

The “Participatory Turn” in Algorithmic Fairness In recent years, interest has grown in applying participatory methods to ML and algorithm design — paralleling a similar historical development in urban planning (Delgado et al. 2023) (see Figure 2). Prior work has advanced the idea that

engaging stakeholders may lead to more equitable algorithmic systems, shift power, and mitigate harm (Suresh et al. 2022; Birhane et al. 2022; Kulynych et al. 2020; Costanza-Chock 2020; D’Ignazio and Klein 2023; Koencke et al. 2023). Recent examples include a case study of QueerInAI (2023), a community-led, grassroots organization, and a participatory algorithm design framework applied to on-demand food donation (Lee et al. 2019). Several other examples are documented in Kim et al. (2025a).

A Pragmatic Turn in Algorithmic Fairness? Recently, there has also been a shift toward pragmatic thinking, such as Green and Viljoen (2020)’s work on “algorithmic realism.” Drawing on the influence of “legal realism” on American legal thinking, they argue that designing fair algorithms must be political, porous, and contextual. Recent calls for a pipeline-aware approach to fairness (Black et al. 2023; Akpinar et al. 2022; Suresh and Gutttag 2021) can also be viewed from a pragmatic perspective; they advocate for an iterative process and are concerned with how bias arises in practice. Lastly, Watson, Mökander, and Floridi (2024) propose “sociotechnical pragmatism” as a compromise to resolve tensions between AI dogmatism and skepticism.

5.3 The Limits of Current Fairness Approaches in Practice

The shift toward participatory and pragmatic methods reflects a growing recognition that current approaches to algorithmic fairness often break down in practice. Holstein et al. (2019) note that practitioners’ needs are at odds with the priorities of most fairness research. The authors call attention to the fact that while much of the algorithmic fairness literature focuses on models, practitioners think primarily about data collection. Deng et al. (2022) and Lee and Singh (2021) both explore how industry practitioners use fairness toolkits, observing a need for more contextualization, integration with existing workflows, and accessible communication with non-technical audiences.

There are many more challenges that practitioners face (Veale, Van Kleek, and Binns 2018; Veale and Binns 2017; Madaio et al. 2020, 2022; Law et al. 2020; Richardson et al. 2021; Gosciak et al. 2026), even when considering holistic fairness approaches. Current fairness efforts are often ad-hoc and limited by organizational constraints (Yan et al. 2026). Some papers offer practical design suggestions: semi-automated bias detection tools with practitioner input, fair “checklists” (Madaio et al. 2020, 2024), or collaborative, experiential knowledge bases (Veale and Binns 2017). The need for these solutions further illustrates the tension between theory and practice.

Furthermore, concerns related to the *mis*-use of participatory methods have arisen. Scholars have raised well-founded critiques that participation in ML and AI is largely “consultative,” narrow, and indifferent to power dynamics (Delgado et al. 2023; Feffer et al. 2023; Birhane et al. 2022; Sloane et al. 2022; Corbett, Denton, and Erete 2023; Groves et al. 2023). Compounding this criticism is the fact that most recent work documenting participatory processes focuses on context-specific, in-depth case studies, from which it can be




Figure 2: The influence of pragmatic ideas on planning theory since Rittel and Webber (1973)’s “wicked problems” (represented by the arrow). Based on Hoch (1984); Healey (2008); Forester (2017). While AI ethics has analogously taken a “participatory turn” (Delgado et al. 2023), reflective and deliberative practice and critical pragmatism have yet to be applied widely.

hard to generalize (Suresh et al. 2024). There is little cohesive and practical guidance about how to structure participatory processes, particularly amidst deep-seated value conflicts. While many note the importance of participation and involving stakeholders, a greater emphasis on listening, deliberation, and modes of communication can make existing participatory approaches more effective.

We turn to urban planning to fill the gap in pragmatic approaches to wicked problems in algorithmic fairness. Urban planning is a field that has championed participation and communication. Drawing on critical pragmatism, we develop flexible recommendations that data science practitioners and organizations can apply to their ongoing work.

6 Urban Planning Tools for AI Ethics

The urban planning shift away from rational-comprehensive planning toward pragmatism, as described in Section 3.1, reflects the tensions that arise between idealized and holistic fairness definitions. In Section 5.2, we discussed how prior work in AI ethics has begun to advocate for participatory and pragmatic approaches, but in Section 5.3, we noted the shortcomings of these methods and ongoing challenges that practitioners face. We can apply learnings from urban planning in addressing similar shortcomings: specifically, regarding pragmatic ideas—and critical pragmatism in particular—as described in Sections 3.1 and 4.

Critical pragmatism provides a set of problem-solving methods that help practitioners answer the question: What can we actually *do*? Here, drawn from the core tenets of critical pragmatism, we propose four recommendations.

1. **Reflect-in-Action:** Practitioners should approach algorithmic fairness like improvisation, iteratively and creatively. This is the essence of the pragmatic tradition. Instead of seeking out fixed “normative truths” about what constitutes ethical and fair algorithms, practitioners should remain open to “intelligently updating” their beliefs as they experiment with different approaches (Anderson 2020). Capturing this dynamic, Schön (2016)’s reflection-in-action is a useful heuristic. It offers a way to manage the context-dependent nature of many algorithmic fairness problems, by approaching the set of unique problems as a professional “repertoire” (Schön 2016). Individuals can reflect-in-action, even when there is no institutional support. Notably, while pragmatism rejects the

rigidity of formal models that function as absolute truths, it appreciates formalisms (e.g., equations, mathematical models) as useful for experimentation and inquiry.

2. **Practice deliberation through public meetings:** Urban planning is known for long public meetings and community engagement processes, but this kind of deliberation – mediating conversations among divided groups – is central for resolving conflict. In a similar vein, participatory methods are becoming more common in ML and AI, but these methods are not necessarily deliberative nor valued as such. Scholars in both planning and AI ethics have cautioned against narrow and extractive forms of participation (Arnstein 1969; Monno and Khakee 2012; Sloane et al. 2022; Delgado et al. 2023; Corbett, Denton, and Erete 2023). Even more justice-oriented approaches often resort to interviews as opposed to sustained public meetings involving real dialogue. We argue that deliberation is not just important for AI governance (Buhmann and Fieseler 2023), but for the entire AI development process — from framing the problem appropriately to safely deploying a tool. Forester (1999)’s *The Deliberative Practitioner* offers practical guidance about how to structure deliberative processes — e.g., framing public meetings as discussions rather than debates, and encouraging informal participatory rituals that build trust (Forester 1999). These meetings are not meant to function like one-sided presentations. They are more similar to community town halls or other political spaces that encourage active participation along with bi-directional communication and learning. They can involve large public discussions or small groups, hundreds of regular meetings or only one.
3. **Listen critically and seek out creative negotiations:** Attend to the role of communication in building trust and resolving conflict. This requires an active listening role, what Forester calls “hermeneutic praxis” (Forester 2013) — a critical interpretive stance to communication that involves understanding both what is said and not said (e.g., the meaning, emotion, and connotation hidden within specific word choices). Critical pragmatists would question stakeholders’ claims to legitimacy — e.g., claims of expertise as a “professional” or over-confident promises with statements like “trust me” (Forester 2013). They would consider how power and status may affect communication: bias a story, distort information, or advance

a specific viewpoint (Forester 1982). These skills are not often valued in technical fields and are more common in qualitative research (Josselson 2004). Creativity is also important, as planners seek to resolve conflict in seemingly impossible settings. One should not “presume impossibility,” and instead attempt to shift power dynamics by anticipating information distortions or through mediation (Forester 1982, 2013).

4. **Share practice-oriented stories:** Forester (1999) argues that practice-oriented storytelling among friends and colleagues, is an important part of learning. These stories are not trivial, nor are they traditional interviews or case studies. They are lengthy, minimally edited, first-hand accounts directly from practitioners about their experiences. They are “rich, morally thick, politically engaged, and organizationally practical” (Forester 1999). Much of Forester’s own work is based on insights derived from practice stories. In fact, understanding his theoretical contributions is most useful in the context of specific examples, which he analyzes over several books (Forester 1999, 1989, 1993). Forester provides a starting point for collecting these stories – e.g., asking “How” rather than “Why” questions or “actor-focused” rather than “spectator-focused” questions (Forester 2012). Practice stories view practitioners as social actors, confronting the messy reality of their work (Wagenaar 2011).

Our recommendations touch on existing calls to action in AI ethics: increase participation, critique power structures, and account for context (Suresh et al. 2024; Costanza-Chock 2020; D’Ignazio and Klein 2023; Koenecke et al. 2023; Kim et al. 2025b). Going a step further, our recommendations are aimed at operationalizing and generalizing from these efforts. Beyond simply consulting stakeholders, we emphasize dialogue, listening, and concern for different modes of communication. These are critical considerations in urban planning that have not been fully utilized in AI ethics nor formalized as specific calls to action. Concepts like reflection-in-action and practice-oriented stories complement these recommendations by helping practitioners learn from and generalize across varied, context-specific tasks.

7 Applying Critical Pragmatism to Algorithmic Fairness: Case Studies

In this section, we envision what a critical pragmatism approach for AI ethics might look like in three case studies. For each case study, we discuss the current approach and how critical pragmatism might lead to improvements. We also identify instances where the framing of critical pragmatism makes explicit strategies that worked well and where critical pragmatism might fall short. We further connect our recommendations to specific points in the algorithm development pipeline in the appendix. We believe our recommendations overall can help with generalizing from these case studies to other settings.

Context varies across each of our case studies. In the first case study, ethical concerns have led to algorithm changes only recently. In the second case study, the algorithm was motivated by concerns related to equity and efficiency. The

third case study is grounded in data feminism, which embeds a wide range of ethical concerns holistically in the project. These three case studies also represent differences in geographic scale. We believe our recommendations apply regardless of scale, whether national or local.

7.1 Case Study 1: Mortgage Lending

The Problem. Given the history of housing inequality in the U.S., fairness and discrimination are paramount concerns for the government regarding mortgages. The Equal Credit Opportunity Act (ECOA) and the Fair Housing Act explicitly forbid discrimination in lending against many sensitive attributes (such as race, gender, national origin), and through the Home Mortgage Disclosure Act, financial institutions are required to disclose data on lending decisions to ensure the ECOA has been upheld (Akinwumi et al. 2021). However, even though explicit lending discrimination is rare in the present-day, prior work has shown that mortgage lending decisions and access to credit tend to be racially skewed (Martinez and Kirchner 2021; Mehrotra et al. 2024).

Current approach. The use of algorithms in mortgage lending decisions is not new. A 2019 report from *The Markup* demonstrated how opaque automated underwriting practices and outdated credit scoring algorithms like “Classic FICO” have led to racially biased outcomes (Martinez and Kirchner 2021). Furthermore, they show how racially disparate outcomes persist even when controlling for many kinds of public data that could potentially explain such disparities. While technical methods have shown promise at reducing inequities, even with biased historical data (Merrill et al. 2024; Singh et al. 2022; Lee and Floridi 2021; Nguyen 2025; Coston, Rambachan, and Chouldechova 2021), such methods are largely irrelevant, as a key impediment to updating the algorithm is political. Until recently, federal agencies have resisted updating “Classic FICO,” even though simple changes (e.g., including rental payments or cashflow in credit-scoring models) could increase credit access to more borrowers (Martinez and Kirchner 2021; Imani Moise 2025; Kumar, Hines, and Dickerson 2022).

Context. Relevant stakeholders include: U.S. federal government agencies, financial institutions, and nonprofit advocates. Financial institutions aim to reduce risk, nonprofit advocates advance equity concerns, and the federal government aims to serve the public interest. While private companies and nonprofits can develop fairer models on their own (Nguyen 2025), many lenders may still opt to follow federal policy guidelines to ensure funding guarantees from the government. Technical fairness methods, including recent proposals from academics, may have little effect without concrete changes in federal policy. While large-scale public engagement processes may be challenging to coordinate at the national level, the federal government is best positioned to oversee such processes and can make far-reaching policy changes.

Alignment with Holistic Fairness: Low. Current practices in mortgage lending, described above, do not reflect a critical pragmatism approach. When considering academic contributions and technical mitigation strategies, most fairness interventions are focused on improving prediction outcomes

and evaluation metrics — not the process of developing the tools themselves. Many of these interventions have been primarily speculative and theoretical.

Applying critical pragmatism: Critical pragmatism would involve a more political and deliberative process through public meetings. The federal government’s automated underwriting process, which is both opaque and largely one-sided, has not supported this kind of process historically (Martinez and Kirchner 2021). Government agencies like the Consumer Financial Protection Bureau (CFPB) and Federal Housing Finance Agency (FHFA) do solicit public feedback, e.g.: public comment on any newly proposed rule, requests for information, and listening sessions (Consumer Financial Protection Bureau 2024; Federal Housing Finance Agency 2022). For written comments, the government evaluates them and decides whether to modify, withdraw, or proceed with each respective rule (Regulations.gov 2024). Requesting public comments results in detached, transactional interactions, not the collaborative and transformational deliberation of critical pragmatism where the actual process of negotiation and communication is important.

The FHFA’s “stakeholder forums and listening sessions” are a promising direction, but these largely include testimony from established industry representatives (Federal Housing Finance Agency 2023b,a, 2022). A critical pragmatism approach would seek to mitigate disparities in participation (reducing material or psychological barriers to participation) and would consider how power may shape communication. By building trust gradually through open discussion, the hope is that a shared set of priorities may emerge. Adopting this kind of approach might have led stakeholders to accept simple methods to improve fairness sooner (e.g., including non-traditional inputs to credit scores (Kumar, Hines, and Dickerson 2022; Imani Moise 2025)).

7.2 Case Study 2: School Choice

The Problem. School choice decisions in many urban areas are complex. NYC — the largest city in the U.S. — assigns students to middle and high-schools using the deferred acceptance (DA) algorithm. The algorithm matches students to schools using students’ preferences and by ranking students based on academic performance, location, and whether the student receives free or reduced-price lunch. This algorithm is widely credited with improving the high school assignment process — both reducing the number of unmatched students and increasing the number of students receiving their first-choice match (Elsen-Rooney 2024b). Prior to 2003, the matching system was inefficient; almost a third of all students did not receive placements to schools they had initially selected (Abdulkadiroğlu, Pathak, and Roth 2005).

Current approach. The DA algorithm has led to significant improvements. The new system is more transparent and the DA algorithm is strategy-proof, meaning it is not possible for students to strategically manipulate outcomes. Still, students may feel unfairly matched, especially if they do not understand the DA algorithm or fail to match to their top schools. News reports highlight the system’s informational complexity and the burden on parents, which may impact students with fewer resources and more limited support sys-

tems (Elsen-Rooney 2024b). To counteract these effects, the city has explored informational and personalized interventions, which empirical evidence suggests are successful (Cohodes et al. 2025; Peng et al. 2026). But, NYC’s school matching process continues to reinforce existing socioeconomic and racial inequalities, particularly through selective admissions screens (Elsen-Rooney 2024a).

A similar problem occurred in San Francisco where the city voted to abandon a school choice algorithm after finding that it exacerbated inequality. Robertson, Nguyen, and Salehi (2021) examined the design and implementation of the algorithm using a Value Sensitive Design approach. They observed several misalignments between modeling assumptions in the algorithm and real-world challenges. Many of these — such as inequitable gaps in access to information or issues of trust — could be mitigated if the design of the algorithm had adopted a critical pragmatism approach.

Context. This case study primarily involves public-sector employees and parents. While public-sector employees may care about broader equity concerns (which motivated the adoption of the DA algorithm in the first place), parents are likely more invested in their children’s welfare. An effective critical pragmatism approach would recognize these sometimes-competing priorities and seek resolution. The fact that the school district manages the DA algorithm provides legitimacy and support for carrying out several of our recommendations, such as organizing public meetings. The local scale can also facilitate deeper engagement with stakeholders, though notably NYC’s school district serves roughly 900,000 students (NYC Public Schools 2026).

Alignment with Holistic Fairness: Moderate. A strength of school choice algorithms like DA is the ability to easily simulate outcomes under a range of conditions. Simulation, much like the sketches that Forester (1999) discusses, can be a tool for the iterative and experimental reflection-in-action process. Reflection-in-action is a method that practitioners can use in any setting, when approaching a new problem or task. A weakness of current approaches to school choice — exemplified in both NYC and San Francisco — is the lack of any deliberative or in-depth public engagement process. One of the main limitations that researchers have cited in prior work is a mismatch in information: misinformed perceptions from parents about how school choice algorithms work and inadequate guidance for students (Cohodes et al. 2025).

Applying critical pragmatism: Deliberation and public meetings could address this problem. These might include interactive workshops and open listening sessions where those responsible for designing and implementing the algorithm are forced to confront the nuanced and conflicting values of different stakeholders. The format of these meetings is critical and importantly helps to develop trust, which has been cited as a significant issue in San Francisco (Robertson, Nguyen, and Salehi 2021, 2022).

Critical pragmatism complements existing methods like value sensitive design (Zhu et al. 2018). Importantly, a critical pragmatism perspective would consider how power and resource differences shape participation (e.g., distorting information, eroding trust, or influencing the opinions stakeholders publicly express (Forester 1989)). In response, crit-

ical pragmatists would adopt an interpretive, “hermeneutic” stance to communication. This stance can help with the broader goal of resolving deep-seated value conflicts, particularly when there are pronounced power imbalances and issues of trust. For example, in the school choice setting, differences in how parents perceive fairness is intertwined with race and socioeconomic status (Nguyen et al. 2024).

A real example of a more deliberative process is an urban planning initiative for NYC School District 15 (D15). The D15 Diversity Plan has been an effort to integrate one of the city’s most socioeconomically and racially segregated school districts prior to 2017 (WXY Studio 2018). It built on years of advocacy, established a multi-stakeholder working group, and convened interactive workshops for families and educators – providing childcare, food, transit passes, and translation to reduce barriers for attending (WXY Studio 2018). The D15 Diversity plan resulted in the removal of all selective screens and a choice-based lottery system with priority for low-income students. As of the 2022-2023 academic school year, District 15 has transformed from the second most segregated school district in the city to the 19th (out of 32) (Closson 2024; WXY Studio 2024).

7.3 Case Study 3: Femicide Counterdata Collection

The Problem. We now turn to counterdata collection for monitoring femicide. Counterdata is a collective, activist response to produce data in settings where powerful institutions (e.g., governments or corporations) fail to do so (D’Ignazio and Klein 2023). A collaboration across both activist and academic organizations, this project used intersectional feminist principles and participatory methodologies to co-design an ML system that identifies relevant news reports of femicide for activists. The project started through interviews with representatives from both activist and civil society organizations (D’Ignazio et al. 2022). Through these interviews, researchers identified the potential for ML to reduce the burden on activists, who regularly sort through violent and upsetting content, by partially automating the detection of femicide cases in news articles (D’Ignazio et al. 2022; Suresh et al. 2022).

Current approach. The femicide counterdata collection project was developed in close alignment with D’Ignazio and Klein (2023)’s work on data feminism. As a result, it demonstrates many of the same values that underlie a critical pragmatism approach. For example, challenging power is one of the core principles of data feminism, and central to this project. The researchers specifically worked with activist groups, rather than well-resourced government agencies or NGOs. They sought to develop a model with high performance on intersectional femicide cases, not only cases for majority groups (Suresh et al. 2022). The project also built on participatory methodologies through a collaborative co-design process that viewed “participation as justice” (Suresh et al. 2022; Sloane et al. 2022).

Context. This case study involves activists and academics, who appear relatively aligned in values and views. This dynamic could be the product of the participatory co-design approach and also the intentional decision to prior-

itize partnering with activists, as opposed to other institutional actors. Though the project does involve participation from stakeholders, the emphasis on trust and deliberation may not be as crucial as in other case studies because there are few high-stakes instances of conflict. The project operates with few regulatory constraints, and the resulting output is closely aligned with the priorities of the activist partners.

Alignment with Holistic Fairness: High. This project reflects high alignment with holistic fairness approaches, as it is based on data feminism (D’Ignazio and Klein 2023). As Suresh et al. (2022) discuss, the co-design ML process was iterative and collaborative, involving stakeholder input through regular evaluation (D’Ignazio 2022). It was also sensitive to context and critical of power. Given this background, critical pragmatism helps clarify what worked well.

Applying critical pragmatism: In particular, this project exemplifies the value in our third recommendation, “creative negotiations” (Forester 2013). Such negotiations are only possible through carefully listening to stakeholders. Given both the desire to memorialize and honor the individuals in the data, and to avoid rigid definitions of femicide, the research team could have decided initially to forego using ML. However, by listening to feedback from activists, they negotiated a more creative, alternative approach involving partial automation: an ML system that generates email alerts to reduce the emotional burden on activists, while activists still retain control of defining the relevance of specific femicide cases. This negotiated outcome was the result of listening to stakeholders, interpreting the underlying meaning in what they shared, and using this interpretation to creatively identify points of overlap.

The femicide counterdata collection project is also useful in demonstrating the trade-offs necessarily arising from a critical pragmatism approach and how it may fall short in some contexts. The research team in this case prioritized working with activists and civil society organizations as a more radical stance in challenging power structures whereas critical pragmatism might have taken a more moderate approach through compromise and negotiation (e.g., engaging with news organizations or government directly).

8 Conclusion

In this paper, we characterize algorithmic fairness problems as “wicked problems,” which are typically the domain of policymakers and urban planners. This characterization suggests a broader analogy between the two fields. We argue that urban planning theory, specifically critical pragmatism, is useful to the AI ethics community. Critical pragmatism focuses on what works *in practice*, addressing gaps in the AI ethics literature between theory and practice. It is grounded in real stories from urban planners. Drawing on critical pragmatism, we offer four recommendations: (1) reflect-in-action, (2) practice deliberation through public meetings, (3) listen critically and negotiate creatively, and (4) share practice-oriented stories. These recommendations underscore the importance of skills often not emphasized in technical, data science settings, but which are crucial for resolving fairness and ethics concerns – skills related to reflexivity, listening, and dialogue.

Acknowledgments

We are grateful to numerous individuals who have provided invaluable feedback and support over the past few years. First, we appreciate John Forester for engaging with our interdisciplinary attempts to bridge two distinct fields, and who consistently replied to our many questions about critical pragmatism. Second, we are grateful to colleagues and faculty at Cornell, who reviewed this work and provided comments – including Angelina Wang, Anthony Townsend, Emma Harvey, Gabriel Agostini, Isabel Corpus, Jae June Lee, Matt Franchi, and Travis Lloyd.

References

- Abdulkadiroğlu, A.; Pathak, P. A.; and Roth, A. E. 2005. The New York City High School Match. *American Economic Review*, 95(2): 364–367.
- Akinwumi, M.; Merrill, J.; Rice, L.; Saleh, K.; and Yap, M. 2021. An AI fair lending policy agenda for the federal financial regulators.
- Akpinar, N.-J.; Nagireddy, M.; Stapleton, L.; Cheng, H.-F.; Zhu, H.; Wu, S.; and Heidari, H. 2022. A Sandbox Tool to Bias(Stress)-Test Fairness Algorithms. arXiv:2204.10233.
- Alexander, E. R. 1981. If planning isn't everything, maybe its something. *The Town Planning Review*, 52(2): 131–142.
- Alexander, E. R. 2022. On planning, planning theories, and practices: A critical reflection. *Planning Theory*, 21(2): 181–211.
- Allmendinger, P.; and Tewdwr-Jones, M. 2002. The communicative turn in urban planning: Unravelling paradigmatic, imperialistic and moralistic dimensions. *Space and polity*, 6(1): 5–24.
- American Planning Association. 2025. Planning Specializations. <https://www.planning.org/choosingplanning/specializations/>. Accessed: 2026-08-07.
- Anderson, E. 2020. How to be a pragmatist. In Chang, R.; and Sylvan, K., eds., *The Routledge Handbook of Practical Reason*, 83–94. Oxford: Routledge.
- Anderson, M. 1964. *The Federal Bulldozer: A Critical Analysis of Urban Renewal, 1949-1962*. Publications of the Joint Center for Urban Studies of the Massachusetts Institute of Technology and Harvard University. Cambridge, Massachusetts: The M.I.T. Press.
- Arif Khan, F.; Manis, E.; and Stoyanovich, J. 2022. Towards Substantive Conceptions of Algorithmic Fairness: Normative Guidance from Equal Opportunity Doctrines. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394772.
- Arnstein, S. R. 1969. A Ladder Of Citizen Participation. *Journal of the American Institute of Planners*, 35(4): 216–224.
- Bacon, M. 2012. *Pragmatism: an introduction*. Cambridge, U.K.: Polity Press.
- Bailey, P. 2022. Addressing the Affordable Housing Crisis Requires Expanding Rental Assistance and Adding Housing Units. <https://www.cbpp.org/research/housing/addressing-the-affordable-housing-crisis-requires-expanding-rental-assistance-and>. Accessed: 2026-08-07.
- Barocas, S.; Hardt, M.; and Narayanan, A. 2023. *Fairness and Machine Learning: Limitations and Opportunities*. The MIT Press.
- Binns, R. 2018. Fairness in Machine Learning: Lessons from Political Philosophy. In Friedler, S. A.; and Wilson, C., eds., *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, 149–159. PMLR.
- Birhane, A.; Isaac, W.; Prabhakaran, V.; Diaz, M.; Elish, M. C.; Gabriel, I.; and Mohamed, S. 2022. Power to the People? Opportunities and Challenges for Participatory AI.
- Black, E.; Naidu, R.; Ghani, R.; Rodolfa, K.; Ho, D.; and Heidari, H. 2023. Toward Operationalizing Pipeline-aware ML Fairness: A Research Agenda for Developing Practical Guidelines and Tools. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '23, 1–11. New York, NY, USA: Association for Computing Machinery. ISBN 9798400703812.
- Buhmann, A.; and Fieseler, C. 2023. Deep Learning Meets Deep Democracy: Deliberative Governance and Responsible Innovation in Artificial Intelligence. *Business Ethics Quarterly*, 33(1): 146–179.
- Callahan, G.; and Ikeda, S. 2014. Jane Jacobs' critique of rationalism in urban planning. *Cosmos + Taxis*, 1(3): 10–19.
- Chohlas-Wood, A.; Coots, M.; Goel, S.; and Nyarko, J. 2023. Designing equitable algorithms. *Nature Computational Science*, 3(7): 601–610.
- Chouldechova, A. 2017. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data*, 5(2): 153–163. PMID: 28632438.
- Churchman, C. W.; Protzen, J.-P.; Webber, M. M.; and Krogh, D. 2007. In Memoriam: Horst W.J. Rittel. *Design Issues*, 23(1): 89–91.
- Closson, T. 2024. What Happened When Brooklyn Tried to Integrate Its Middle Schools. <https://www.nytimes.com/2024/06/20/nyregion/brooklyn-middle-schools-integration.html>. Accessed: 2026-08-07.
- Cohodes, S. R.; Corcoran, S. P.; Jennings, J. L.; and Sattin-Bajaj, C. 2025. When do informational interventions work? Experimental evidence from New York City high school choice. *Educational Evaluation and Policy Analysis*, 47(1): 208–236.
- Colin Lecher, M. V. 2021. NYC's School Algorithms Cement Segregation. This Data Shows How. <https://themarkup.org/machine-learning/2021/05/26/nycs-school-algorithms-cement-segregation-this-data-shows-how>. Accessed: 2026-08-07.
- Consumer Financial Protection Bureau. 2024. Notice and opportunities to comment. <https://www.consumerfinance>.

gov/rules-policy/notice-opportunities-comment/. Accessed: 2024-01-24.

Corbett, E.; Denton, R.; and Erete, S. 2023. Power and Public Participation in AI. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '23. New York, NY, USA: Association for Computing Machinery. ISBN 9798400703812.

Costanza-Chock, S. 2020. *Design justice: Community-led practices to build the worlds we need*. Cambridge, Massachusetts: The MIT Press.

Coston, A.; Rambachan, A.; and Chouldechova, A. 2021. Characterizing Fairness Over the Set of Good Models Under Selective Labels. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 2144–2155. PMLR.

Davidoff, P. 1965. Advocacy and pluralism in planning. *Journal of the American Institute of planners*, 31(4): 331–338.

Davis, J. 2021. The double-edged sword of upzoning. <https://www.brookings.edu/articles/the-double-edged-sword-of-upzoning/>. Accessed: 2026-08-07.

Davis, J. L.; Williams, A.; and Yang, M. W. 2021. Algorithmic reparation. *Big Data & Society*, 8(2): 20539517211044808.

Delgado, F.; Yang, S.; Madaio, M.; and Yang, Q. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '23. New York, NY, USA: Association for Computing Machinery. ISBN 9798400703812.

Deng, W. H.; Barocas, S.; and Wortman Vaughan, J. 2025. Supporting Industry Computing Researchers in Assessing, Articulating, and Addressing the Potential Negative Societal Impact of Their Work. *Proc. ACM Hum.-Comput. Interact.*, 9(2).

Deng, W. H.; Nagireddy, M.; Lee, M. S. A.; Singh, J.; Wu, Z. S.; Holstein, K.; and Zhu, H. 2022. Exploring How Machine Learning Practitioners (Try To) Use Fairness Toolkits. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 473–484. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393522.

Desmond, M.; and Gershenson, C. 2016. Housing and employment insecurity among the working poor. *Social problems*, 63(1): 46–67.

Dewey, J. 1927. *The Public and Its Problems*. Denver: Alan Swallow. <https://www.gutenberg.org/cache/epub/71000/pg71000-images.html>. Accessed: 2026-08-07.

D'Ignazio, C. 2022. *Counting Feminicide: Data Feminism in Action*. The MIT Press. <https://mitpressonpubpub.mitpress.mit.edu/counting-feminicide>. Accessed: 2026-08-07.

D'Ignazio, C.; Cruxên, I.; Val, H. S.; Cuba, A. M.; García-Montes, M.; Fumega, S.; Suresh, H.; and So, W. 2022. Feminicide and counterdata production: Activist efforts to monitor and challenge gender-related violence. *Patterns*, 3(7).

D'Ignazio, C.; and Klein, L. F. 2023. *Data feminism*. Cambridge, Massachusetts: The MIT press.

Dobbe, R.; Krendl Gilbert, T.; and Mintz, Y. 2021. Hard choices in artificial intelligence. *Artificial Intelligence*, 300: 103555.

Elsen-Rooney, M. 2024a. How the NYC high school admissions process sorts kids by race, poverty, disability. <https://www.chalkbeat.org/newyork/2024/10/16/nyc-high-school-admissions-sorts-students-by-race-poverty-and-disability/>. Accessed: 2026-08-07.

Elsen-Rooney, M. 2024b. NYC reformed high school admissions 20 years ago. Did it make things better? <https://www.chalkbeat.org/newyork/2024/12/16/high-school-admissions-reforms-segregation-inequality-conference/>. Accessed: 2026-08-07.

Fabris, A.; Baranowska, N.; Dennis, M. J.; Graus, D.; Hacker, P.; Saldivar, J.; Zuiderveen Borgesius, F.; and Biega, A. J. 2025. Fairness and Bias in Algorithmic Hiring: A Multidisciplinary Survey. *ACM Trans. Intell. Syst. Technol.*, 16(1).

Fazelpour, S.; and Lipton, Z. C. 2020. Algorithmic Fairness from a Non-ideal Perspective. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, AIES '20, 57–63. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-7110-0.

Federal Housing Finance Agency. 2022. Public Listening Session on Credit Score Models. <https://www.fhfa.gov/news/videos/public-listening-session-on-credit-score-models>. Accessed: 2026-08-07.

Federal Housing Finance Agency. 2023a. FHFA Announces Next Phase of Public Engagement Process for Updated Credit Score Requirements. <https://www.fhfa.gov/Media/PublicAffairs/Pages/FHFA-Announces-Next-Phase-of-Public-Engagement-Process-for-Updated-Credit-Score-Requirements.aspx>. Accessed: 2025-01-21.

Federal Housing Finance Agency. 2023b. Provide Input. <https://www.fhfa.gov/AboutUs/Contact/Pages/Request-for-Information-Form.aspx>. Accessed: 2024-01-21.

Feffer, M.; Skirpan, M.; Lipton, Z.; and Heidari, H. 2023. From Preference Elicitation to Participatory ML: A Critical Survey & Guidelines for Future Research. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 38–48. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702310.

Fischer, W.; Acosta, S.; and Bailey, A. 2021. An Agenda for the Future of Public Housing. <https://www.cbpp.org/research/housing/an-agenda-for-the-future-of-public-housing>. Accessed: 2024-01-21.

Forester, J. 1980. Critical theory and planning practice. *Journal of the American Planning Association*, 46(3): 275–286.

- Forester, J. 1981. Questioning and organizing attention: Toward a critical theory of planning and administrative practice. *Administration & Society*, 13(2): 161–205.
- Forester, J. 1982. Planning in the Face of Power. *Journal of the American Planning Association*, 48(1): 67–80.
- Forester, J. 1999. *The deliberative practitioner: Encouraging participatory planning processes*. Cambridge, Massachusetts: The MIT Press.
- Forester, J. 2012. Learning to improve practice: Lessons from practice stories and practitioners' own discourse analyses (or why only the loons show up). *Planning Theory & Practice*, 13(1): 11–26.
- Forester, J. 2013. On the theory and practice of critical pragmatism: Deliberative practice and creative negotiations. *Planning Theory*, 12(1): 5–22. Publisher: Sage Publications, Ltd.
- Forester, J. 2017. In Haselsberger, B., ed., *Encounters in Planning Thought*, chapter On the evolution of a critical pragmatism, 280–296. New York: Routledge.
- Forester, J. 2020. Five generations of theory–practice tensions: enriching socio-ecological practice research. *Socio-Ecological Practice Research*, 2: 111–119.
- Forester, J. F. 1989. *Planning in the Face of Power*. Berkeley and Los Angeles, California: University of California Press.
- Forester, J. F. 1993. *Critical Theory, Public Policy, and Planning Practice*. Albany, New York: State University of New York Press.
- Freire, P. 1970. *Pedagogy of the Oppressed*. New York, New York: Continuum Books.
- Friedmann, J. 1987. *Planning in the Public Domain: From Knowledge to Action*. Princeton, NJ: Princeton University Press.
- Friedmann, J. 2008. The transactive style of planning. In Hillier, J.; and Healey, P., eds., *Contemporary Movements in Planning Theory: Critical Essays in Planning Theory: Volume 3*, 171–193. Routledge.
- Gans, H. J. 1959. The human implications of current redevelopment and relocation planning. *Journal of the American Institute of Planners*, 25(1): 15–26.
- Gerchick, M. K.; Encarnación, R.; Tanigawa-Lau, C.; Armstrong, L.; Gutiérrez, A.; and Metaxa, D. 2025. Auditing the Audits: Lessons for Algorithmic Accountability from Local Law 144's Bias Audits. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '25, 29–44. New York, NY, USA: Association for Computing Machinery. ISBN 9798400714825.
- Gosciak, J.; Boyce, L.; Wang, A.; and Koenecke, A. 2026. *Scrutinizing Index-Based Risk Assessments: A Case Study in NYC Decision-making for Heat Emergency Management*, 6659–6701. New York, NY, USA: Association for Computing Machinery. ISBN 9798400725968.
- Green, B. 2019. *The smart enough city: Putting technology in its place to reclaim our urban future*. Cambridge, Massachusetts: The MIT Press.
- Green, B. 2020. The false promise of risk assessments: epistemic reform and the limits of fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, 594–606. New York, NY, USA: Association for Computing Machinery. ISBN 9781450369367.
- Green, B. 2022. Escaping the impossibility of fairness: From formal to substantive algorithmic fairness. *Philosophy & Technology*, 35(4): 90.
- Green, B.; and Viljoen, S. 2020. Algorithmic realism: expanding the boundaries of algorithmic thought. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, 19–31. New York, NY, USA: Association for Computing Machinery. ISBN 9781450369367.
- Groves, L.; Peppin, A.; Strait, A.; and Brennan, J. 2023. Going public: the role of public participation approaches in commercial AI labs. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, 1162–1173. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701924.
- Gubits, D.; Shinn, M.; Bell, S.; Wood, M.; Dastrup, S. R.; Solari, C.; Brown, S.; Brown, S.; Dunton, L.; Lin, W.; et al. 2015. Family options study: Short-term impacts of housing and services interventions for homeless families. Technical report, US Department of Housing and Urban Development, Office of Policy Development and Research.
- Habermas, J. 1979. *Communication and the Evolution of Society*. Boston, Massachusetts: Beacon Press.
- Hall, P. 1989. The turbulent eighth decade: Challenges to American city planning. *Journal of the American Planning Association*, 55(3): 275–282.
- Healey, P. 1992. Planning through Debate: The Communicative Turn in Planning Theory. *The Town Planning Review*, 63(2): 143–162. Publisher: Liverpool University Press.
- Healey, P. 2008. The Pragmatic Tradition in Planning Thought. *Journal of Planning Education and Research*. Publisher: SAGE PublicationsSage CA: Los Angeles, CA.
- Heidari, H.; Loi, M.; Gummadi, K. P.; and Krause, A. 2019. A Moral Framework for Understanding Fair ML through Economic Models of Equality of Opportunity. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* '19, 181–190. New York, NY, USA: Association for Computing Machinery. ISBN 9781450361255.
- Hoch, C. J. 1984. Pragmatism, planning, and power. *Journal of Planning Education and Research*, 4(2): 86–95.
- Holstein, K.; Wortman Vaughan, J.; Daumé, H.; Dudik, M.; and Wallach, H. 2019. Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need? In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, 1–16. New York, NY, USA: Association for Computing Machinery. ISBN 9781450359702.
- Hunt, D. B. 2001. What went wrong with public housing in Chicago? A history of the Robert Taylor homes. *Journal of the Illinois State Historical Society (1998-)*, 94(1): 96–123.

- Huxley, M. 2000. The limits to communicative planning. *Journal of planning education and research*, 19(4): 369–377.
- Huxley, M.; and Yiftachel, O. 2000. New paradigm or old myopia? Unsettling the communicative turn in planning theory. *Journal of planning education and research*, 19(4): 333–342.
- Imani Moise, K. H. 2025. Challenger to FICO Credit Scores Gets Green Light for Use in Mortgages. <https://www.wsj.com/finance/regulation/mortgage-loan-vantagesscore-fico-challenge-01841472>. Accessed: 2026-08-07.
- Innes, J. E. 1995. Planning theory’s emerging paradigm: Communicative action and interactive practice. *Journal of planning education and research*, 14(3): 183–189.
- Jacobs, A. Z.; and Wallach, H. 2021. Measurement and Fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’21, 375–385. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383097.
- Jacobs, J. 1992. *The Death and Life of Great American Cities*. Vintage Books. New York: Random House Inc. ISBN 9780679741954.
- Josselson, R. 2004. The hermeneutics of faith and the hermeneutics of suspicion. *Narrative inquiry*, 14(1): 1–28.
- Kalluri, P. 2020. Don’t ask if artificial intelligence is good or fair, ask how it shifts power. *Nature*, 583(7815): 169–169.
- Kim, E.; Balloli, V.; Karimian, B.; Bondi-Kelly, E.; and Fish, B. 2025a. Relationship Building: Analyzing Recruitment Strategies for Participatory AI. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(2): 1426–1439.
- Kim, R. M.; Kuehnert, B.; Lazar, S.; Singh, R.; and Heidari, H. 2025b. The AI Power Disparity Index: Toward a Compound Measure of AI Actors’ Power to Shape the AI Ecosystem. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(2): 1453–1464.
- Kleinberg, J.; Mullainathan, S.; and Raghavan, M. 2016. Inherent Trade-Offs in the Fair Determination of Risk Scores. arXiv:1609.05807.
- Koenecke, A.; Giannella, E.; Willer, R.; and Goel, S. 2023. Popular Support for Balancing Equity and Efficiency in Resource Allocation: A Case Study in Online Advertising to Increase Welfare Program Awareness. volume 17, 494–506.
- Kulynych, B.; Madras, D.; Milli, S.; Raji, I. D.; Zhou, A.; and Zemel, R. 2020. Participatory approaches to machine learning. In *International Conference on Machine Learning Workshop*, volume 7.
- Kumar, I. E.; Hines, K. E.; and Dickerson, J. P. 2022. Equalizing Credit Opportunity in Algorithms: Aligning Algorithmic Fairness Research with U.S. Fair Lending Regulation. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’22, 357–368. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392471.
- Law, P.-M.; Malik, S.; Du, F.; and Sinha, M. 2020. Designing Tools for Semi-Automated Detection of Machine Learning Biases: An Interview Study. arXiv:2003.07680.
- Lee, M. K.; Kusbit, D.; Kahng, A.; Kim, J. T.; Yuan, X.; Chan, A.; See, D.; Noothigattu, R.; Lee, S.; Psomas, A.; and Procaccia, A. D. 2019. WeBuildAI: Participatory Framework for Algorithmic Governance. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW).
- Lee, M. S. A.; and Floridi, L. 2021. Algorithmic fairness in mortgage lending: from absolute conditions to relational trade-offs. *Minds and Machines*, 31(1): 165–191.
- Lee, M. S. A.; Floridi, L.; and Singh, J. 2021. Formalising trade-offs beyond algorithmic fairness: lessons from ethical philosophy and welfare economics. *AI and Ethics*, 1(4): 529–544.
- Lee, M. S. A.; and Singh, J. 2021. The Landscape and Gaps in Open Source Fairness Toolkits. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380966.
- Lindblom, C. E. 1959. The science of “muddling” through. *Public Administration Review*, 19(2): 79–88.
- Lum, K.; and Isaac, W. 2016. To predict and serve? *Significance*, 13(5): 14–19.
- Maclay, K. 2006. Urban planning pioneer Melvin Webber dies at 86. https://newsarchive.berkeley.edu/news/media/releases/2006/12/01_webber.shtml. Accessed: 2026-08-07.
- Madaio, M.; Egede, L.; Subramonyam, H.; Wortman Vaughan, J.; and Wallach, H. 2022. Assessing the Fairness of AI Systems: AI Practitioners’ Processes, Challenges, and Needs for Support. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW1): 52:1–52:26.
- Madaio, M. A.; Chen, J.; Wallach, H.; and Wortman Vaughan, J. 2024. Tinker, Tailor, Configure, Customize: The Articulation Work of Contextualizing an AI Fairness Checklist. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW1): 1–20.
- Madaio, M. A.; Stark, L.; Wortman Vaughan, J.; and Wallach, H. 2020. Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI ’20, 1–14. New York, NY, USA: Association for Computing Machinery. ISBN 9781450367080.
- Maddy Varner, C. L. 2021. How We Investigated NYC High School Admissions. <https://themarkup.org/show-your-work/2021/05/26/how-we-investigated-nyc-high-school-admissions>. Accessed: 2026-08-07.
- Martinez, E.; and Kirchner, L. 2021. The Secret Bias Hidden in Mortgage-Approval Algorithms. <https://themarkup.org/denied/2021/08/25/the-secret-bias-hidden-in-mortgage-approval-algorithms>. Accessed: 2024-01-24.
- Mehrotra, A.; Pang, D.; Zhu, J.; Choi, J. H.; and Ratcliffe, J. 2024. Evidence of Disparities in Access to Mortgage Credit. Technical report, Urban Institute.

- Merrill, J.; Jones, M.; Eberstein, M.; Saleh, K.; Lockwood, D.; Petrosyan, L.; and Akinwum, M. 2024. Improving mortgage underwriting and pricing outcomes for protected classes through distribution matching. *National Fair Housing Alliance and FairPlay*.
- Misak, C. 2004. Making disagreement matter: Pragmatism and deliberative democracy. *The Journal of Speculative Philosophy*, 18(1): 9–22.
- Monno, V.; and Khakee, A. 2012. Tokenism or political activism? Some reflections on participatory planning. *International Planning Studies*, 17(1): 85–101.
- National Low Income Housing Coalition. 2023. Out of Reach. Technical report.
- Nguyen, F. 2025. Training AI to Tackle Bias in the Mortgage Industry. <https://shelterforce.org/2025/06/19/training-ai-to-tackle-bias-in-the-mortgage-industry/>. Accessed: 2026-01-24.
- Nguyen, T.; Alam, S.; Hu, C.; Albiston, C.; and Salehi, N. 2024. Definitions of Fairness Differ Across Socioeconomic Groups & Shape Perceptions of Algorithmic Decisions. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW2): 1–31.
- NYC Public Schools. 2026. NYCPS Data at a Glance. <https://www.schools.nyc.gov/about-us/reports/nycps-data-at-a-glance>. Accessed: 2026-08-07.
- Organizers of QueerInAI; Ovalle, A.; Subramonian, A.; Singh, A.; Voelcker, C.; Sutherland, D. J.; Locatelli, D.; Breznik, E.; Klubicka, F.; Yuan, H.; J, H.; Zhang, H.; Shriram, J.; Lehman, K.; Soldaini, L.; Sap, M.; Deisenroth, M. P.; Pacheco, M. L.; Ryskina, M.; Mundt, M.; Agarwal, M.; Mclean, N.; Xu, P.; Pranav, A.; Korpan, R.; Ray, R.; Mathew, S.; Arora, S.; John, S.; Anand, T.; Agrawal, V.; Agnew, W.; Long, Y.; Wang, Z. J.; Talat, Z.; Ghosh, A.; Dennler, N.; Noseworthy, M.; Jha, S.; Baylor, E.; Joshi, A.; Bilenko, N. Y.; Mcnamara, A.; Gontijo-Lopes, R.; Markham, A.; Dong, E.; Kay, J.; Saraswat, M.; Vytla, N.; and Stark, L. 2023. Queer In AI: A Case Study in Community-Led Participatory AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, 1882–1895. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701924.
- Peng, K.; Ryu, E.; Kleinberg, J.; Tardos, ; and Garg, N. 2026. Connecting Application Behavior to Undermatching in New York City School Choice. Forthcoming.
- Pessach, D.; and Shmueli, E. 2023. Algorithmic fairness. In *Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook*, 867–886. Springer.
- Regulations.gov. 2024. Learn About the Regulatory Process. <https://www.regulations.gov/learn>. Accessed: 2024-01-24.
- Richardson, B.; Garcia-Gathright, J.; Way, S. F.; Thom, J.; and Cramer, H. 2021. Towards Fairness in Practice: A Practitioner-Oriented Rubric for Evaluating Fair ML Toolkits. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380966.
- Rittel, H. W. J.; and Webber, M. M. 1973. Dilemmas in a general theory of planning. *Policy Sciences*, 4(2): 155–169.
- Robertson, S.; Nguyen, T.; and Salehi, N. 2021. Modeling Assumptions Clash with the Real World: Transparency, Equity, and Community Challenges for Student Assignment Algorithms. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380966.
- Robertson, S.; Nguyen, T.; and Salehi, N. 2022. Not Another School Resource Map: Meeting Underserved Families' Information Needs Requires Trusting Relationships and Personalized Care. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2).
- Robinson, D. G. 2022. *Voices in the code: a story about people, their values, and the algorithm they made*. New York: Russell Sage Foundation.
- Scantamburlo, T. 2021. Non-empirical problems in fair machine learning. *Ethics and Information Technology*, 23(4): 703–712.
- Schön, D. A. 2016. *The Reflective Practitioner: How Professionals Think in Action*. London, U.K.: Routledge. ISBN 978-1-315-23747-3.
- Scott, J. C. 1998. *Seeing like a state: How certain schemes to improve the human condition have failed*. New Haven, Connecticut: Yale University Press.
- Selbst, A. D.; Boyd, D.; Friedler, S. A.; Venkatasubramanian, S.; and Vertesi, J. 2019. Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* '19, 59–68. New York, NY, USA: Association for Computing Machinery. ISBN 9781450361255.
- Sendak, M.; Elish, M. C.; Gao, M.; Futoma, J.; Ratliff, W.; Nichols, M.; Bedoya, A.; Balu, S.; and O'Brien, C. 2020. “The human body is a black box”: supporting clinical decision-making with deep learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, 99–109. New York, NY, USA: Association for Computing Machinery. ISBN 9781450369367.
- Shi, Z. R.; Yuan, Y.; Lo, K.; Kapoor, A.; Levin-Decanini, A.; Hudson, S.; Tepperman, J.; Lizarondo, L.; and Fang, F. 2022. AI for Food Rescue. In Tambe, M.; Fang, F.; and Wilder, B., eds., *AI for Social Impact*. <https://ai4sibook.org/>.
- Singh, A.; Singh, J.; Khan, A.; and Gupta, A. 2022. Developing a novel fair-loan classifier through a multi-sensitive debiasing pipeline: Dualfair. *Machine Learning and Knowledge Extraction*, 4(1): 240–253.
- Sloane, M.; Moss, E.; Awomolo, O.; and Forlano, L. 2022. Participation Is not a Design Fix for Machine Learning. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394772.

- Smith, J. J.; Madaio, M.; Burke, R.; and Fiesler, C. 2025. Pragmatic Fairness: Evaluating ML Fairness Within the Constraints of Industry. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '25, 628–638. New York, NY, USA: Association for Computing Machinery. ISBN 9798400714825.
- Strauß, S. 2021. Deep automation bias: How to tackle a wicked problem of ai? *Big Data and Cognitive Computing*, 5(2): 18.
- Suresh, H.; and Guttat, J. 2021. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. In *Proceedings of the 1st ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450385534.
- Suresh, H.; Movva, R.; Dogan, A. L.; Bhargava, R.; Cruxen, I.; Cuba, A. M.; Taurino, G.; So, W.; and D'Ignazio, C. 2022. Towards Intersectional Feminist and Participatory ML: A Case Study in Supporting Femicide Counterdata Collection. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 667–678. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393522.
- Suresh, H.; Tseng, E.; Young, M.; Gray, M.; Pierson, E.; and Levy, K. 2024. Participation in the age of foundation models. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '24, 1609–1621. New York, NY, USA: Association for Computing Machinery. ISBN 9798400704505.
- The New York City Council. 2021. Automated employment decision tools. 2021/144. <https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=4344524&GUID=B051915D-A9AC-451E-81F8-6596032FA3F9>. Accessed: 2026-08-07.
- Veale, M.; and Binns, R. 2017. Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data. *Big Data & Society*, 4(2): 2053951717743530.
- Veale, M.; Van Kleek, M.; and Binns, R. 2018. Fairness and Accountability Design Needs for Algorithmic Support in High-Stakes Public Sector Decision-Making. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, 1–14. New York, NY, USA: Association for Computing Machinery. ISBN 9781450356206.
- Von Hoffman, A. 2008. The lost history of urban renewal. *Journal of Urbanism*, 1(3): 281–301.
- Wagenaar, H. 2011. “A beckon to the makings, workings and doings of human beings”: The critical pragmatism of John Forester. *Public Administration Review*, 71(2): 293–298.
- Wagenaar, H. 2014. *Meaning in action: Interpretation and dialogue in policy analysis*. New York: Routledge.
- Watson, D. S.; Mökander, J.; and Floridi, L. 2024. Competing narratives in AI ethics: a defense of sociotechnical pragmatism. *AI & SOCIETY*, 1–23.
- Weerts, H.; Dudík, M.; Edgar, R.; Jalali, A.; Lutz, R.; and Madaio, M. 2023. Fairlearn: Assessing and Improving Fairness of AI Systems. *Journal of Machine Learning Research*, 24(257): 1–8.
- Whittemore, A. H. 2015. Practitioners Theorize, Too: Reaffirming Planning Theory in a Survey of Practitioners' Theories. *Journal of Planning Education and Research*, 35(1): 76–85. Publisher: SAGE Publications Inc.
- WXY Studio. 2018. D15 Diversity Plan. https://media.wxy-labs.org/190620_D15_Diversity_Plan_Final_Report.e1efc14f94.pdf. Accessed: 2026-08-07.
- WXY Studio. 2024. District 15 Diversity Plan Evaluation & Reflection. https://media.wxy-labs.org/1/District_15_Diversity_Plan_Evaluation_Reflection.14ffab792c.pdf. Accessed: 2026-08-07.
- Yan, J. N.; Harvey, E.; Wang, J.; Rzeszutarski, J. M.; and Koenecke, A. 2026. Fairness-in-the-Workflow: How Machine Learning Practitioners at Big Tech Companies Approach Fairness in Recommender Systems. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI '26. New York, NY, USA: Association for Computing Machinery. ISBN 9798400722783.
- Yiftachel, O.; and Huxley, M. 2000. Debating Dominance and Relevance: Notes on the 'Communicative Turn' in Planning Theory. *International Journal of Urban and Regional Research*, 24(4): 907–913.
- Zhu, H.; Yu, B.; Halfaker, A.; and Terveen, L. 2018. Value-Sensitive Algorithm Design: Method, Case Study, and Lessons. *Proc. ACM Hum.-Comput. Interact.*, 2(CSCW).