

Beyond Solutionism: Techno-Realism as a Design Framework for Responsible Intervention in Human-Centered Innovation

ANONYMOUS AUTHOR(S)

Interactive computing and artificial intelligence are increasingly positioned as solutions to complex social problems, yet dominant problem–solution paradigms in HCI often produce underwhelming or harmful outcomes when technologies are deployed under conditions of structural inequality, institutional constraint, and asymmetric power. This paper introduces the Techno-Realist Design Framework, a practice-grounded approach that reorients design from universalist problem-solving toward conditional intervention, moral constraint, and institutional accountability. The framework is derived through auto-ethnographic synthesis of longitudinal justice-oriented design work across community, civic, and organizational contexts. Techno-Realism articulates seven design commitments paired with evaluative tests that function as procedural checkpoints, producing concrete artifacts for traceability, deliberation, and repair across the technology lifecycle. Rather than prescribing a fixed method, the framework supports designers in deciding not only how to build responsibly, but when refusal, delay, or withdrawal constitute ethically appropriate design outcomes. The paper contributes a reflexive design framework for navigating ethical decision-making under real-world constraints.

CCS Concepts: • **Human-centered computing** → **Collaborative and social computing**; • **Social and professional topics** → *Computing / technology policy*; Codes of ethics.

Additional Key Words and Phrases: Techno-Realism, Human-Centered Computing, Justice-Oriented Design, Techno-Solutionism, Responsible Innovation, Public Interest Technology, Community-Engaged Design

ACM Reference Format:

Anonymous Author(s). 2026. Beyond Solutionism: Techno-Realism as a Design Framework for Responsible Intervention in Human-Centered Innovation. In *Proceedings of ACM SIGCHI Conference on Designing Interactive Systems (ACM DIS '26)*. ACM, New York, NY, USA, 20 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 Introduction

Interactive computing and artificial intelligence are increasingly positioned as solutions to some of society’s most enduring social problems, including inequality, exclusion, inefficiency, and harm. Within human–computer interaction (HCI), this orientation has crystallized into a dominant *problem–solution* paradigm, in which social issues are rendered legible as design challenges and addressed through technical artifacts, interfaces, and data-driven systems [39, 52]. While this paradigm has produced important advances in usability, accessibility, and participatory methods, it has also enabled a persistent form of *techno-solutionism*: the assumption that complex, structural, and historically rooted injustices can be meaningfully addressed through better design, smarter algorithms, or more inclusive datasets alone [14, 37].

A growing body of HCI, DIS, CSCW, and FAccT scholarship has challenged this assumption, demonstrating that technological interventions aimed at social good often produce underwhelming, fragile, or even harmful outcomes [10, 12, 33, 48]. These failures are not merely technical but epistemic and political. By abstracting social problems into

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

Manuscript submitted to ACM

53 discrete design requirements, problem–solution logics frequently detach harm from its structural origins, reframe
54 collective injustice as individual deficit, and privilege institutional or commercial stakeholders over those most affected
55 by technological decision-making [33, 38]. As a result, technologies intended to promote inclusion or equity may instead
56 stabilize existing hierarchies of power, visibility, and control.

57
58 This fixation on technological solutions is deeply entangled with the political economy of contemporary innovation.
59 Capitalization, scale, and speed shape what counts as legitimate progress in technology research and design, incen-
60 tivizing intervention over restraint and deployment over accountability [3, 6]. In these contexts, social problems are
61 reframed as opportunities for innovation, investment, and growth, rather than as symptoms of historical dispossession,
62 racial capitalism, or institutional neglect. HCI scholars have shown how this dynamic produces technologies that
63 are simultaneously framed as benevolent and experienced as extractive, particularly in domains such as surveillance,
64 finance, healthcare, and language technologies [15, 31, 36].

65
66 For communities that have been historically overlooked, marginalized, or excluded, technology is rarely encountered
67 as a neutral tool. Instead, it operates as a mediator of recognition, access, risk, and legitimacy—shaping who is seen,
68 heard, believed, or denied [29, 45]. Black communities in the United States offer a particularly stark illustration of these
69 dynamics due to the long entanglement of technology with racialized surveillance, extraction, and governance [15, 38].
70 However, the underlying conditions are not unique. Similar patterns emerge globally among Indigenous communities,
71 racialized minorities, disabled people, migrants, and economically precarious populations whose lives are increasingly
72 governed by digital infrastructures they neither designed nor control [5, 23]. In such contexts, innovation often arrives
73 not as empowerment but as exposure—to misrecognition, heightened monitoring, or institutional harm.

74
75 Responding to these challenges requires more than refining existing design methods or appending ethical considera-
76 tions at the margins of technical work. It demands a fundamental reorientation of how innovation is conceptualized,
77 justified, and practiced in HCI. This paper introduces the *Techno-Realist Design Framework*, a justice-oriented approach
78 that resists universalist problem-solving and instead centers socially situated realism: an explicit reckoning with
79 history, power, responsibility, and the unequal distribution of technological benefits and harms. Drawing on long-term,
80 community-engaged research and design practice, Techno-Realism foregrounds normative guidance, structural analysis,
81 and accountability as core design commitments. Rather than asking how technology can solve social problems, the
82 framework asks what forms of technological intervention are morally permissible, for whom, under what conditions,
83 and with what obligations over time—including the obligation to refuse or withdraw technologies when they risk
84 compounding harm [43, 50]. In doing so, this framework offers a necessary foundation for justice-oriented innovation
85 that extends beyond any one particular community, yet to all those whose lives are shaped by technologies designed
86 without them.

92 93 **2 Background**

94 95 **2.1 Techno-Solutionism and the Political Economy of Innovation**

96
97 HCI’s dominant orientation toward problem-solving is deeply entangled with the political economy of techno-
98 solutionism, which presumes that technology can and should address social problems at scale. The term *techno-*
99 *solutionism* describes the tendency to treat technology as the primary or only answer to complex social issues, situating
100 computational systems as solutions independent of political, economic, or structural context [14, 37]. Critics argue
101 that this impulse derives not only from disciplinary habit but also from the political economy of innovation, where
102 capitalization, scalability, and disruption are valorized over care, long-term accountability, and contextual sensitivity
103

[3, 6]. Under these conditions, social issues become market opportunities rather than sites for collective deliberation or structural change, with designers incentivized to optimize for novelty and impact metrics rather than justice or sustainability.

Techno-solutionism thus operates as both epistemic assumption and institutional imperative. Even in work framed as socially beneficial, such as responsible innovation or participatory approaches, the underlying logic often remains fixated on intervention through technology rather than critical engagement with the conditions that produce harm in the first place. This dynamic has been especially evident in domains such as predictive algorithms in policing, algorithmic decision-making in credit and employment, and automated language technologies, where solutions often perpetuate or deepen existing inequities despite intentions to promote fairness or inclusion [31, 38]. In this context, critical scholars have called for approaches that explicitly interrogate the limits and harms of technological intervention, rather than assuming the primacy of solutions themselves.

2.2 Techno-Optimism, Techno-Pessimism, and the Case for Techno-Realism

Contemporary discourse on technological innovation is often polarized between techno-optimist and techno-pessimist orientations, each of which carries distinct assumptions about the role of computing in social change. Techno-optimist perspectives position technological innovation as an inherently progressive force capable of scaling solutions to complex social problems through efficiency, automation, and disruption. This orientation is deeply intertwined with venture capital logics, platform economies, and engineering cultures that privilege rapid deployment, growth, and abstraction [3, 35, 54]. Within HCI and adjacent fields, techno-optimism frequently manifests as solutionism: the tendency to frame social and political problems as amenable to technical fixes, often detached from their historical and structural conditions [14, 37]. Prior work has demonstrated how such approaches routinely produce underwhelming, fragile, or harmful outcomes when deployed in contexts shaped by inequality, racialization, and institutional power [10, 12, 18].

Closely related to techno-optimism are concepts of techno-determinism and techno-fixation. Techno-determinism treats technological development as an autonomous and inevitable force that shapes social arrangements independent of human agency or political choice [58]. Techno-fixation, in turn, describes the persistent belief that technological intervention is the primary or default pathway for addressing social problems, even when evidence suggests otherwise [32]. Together, these orientations reinforce a narrow imagination of design agency in which innovation is framed as both necessary and morally justified by default, limiting space for refusal, restraint, or institutional accountability. While analytically distinct, techno-determinism and techno-fixation often function as enabling logics for techno-optimism by naturalizing technological inevitability and narrowing the space for political choice.

In contrast, techno-pessimist perspectives emphasize the structural harms, extractive political economies, and colonial or racialized histories embedded in technological systems. Critical scholarship across HCI, STS, and media studies has documented how computing infrastructures reproduce surveillance, dispossession, labor exploitation, and epistemic violence [11, 15, 38]. Within design research, speculative and critical approaches have sometimes foregrounded refusal, non-use, and anti-instrumentality as strategies for resisting technological domination [9, 21]. While these critiques have been vital for exposing systemic harm, purely pessimistic orientations can risk political paralysis or disengagement from the institutional sites where technological power is actively exercised [24]. Refusal alone does not resolve the governance challenges posed by systems already embedded in public infrastructure, financial institutions, healthcare, and state administration.

In response to these polarized orientations, Techno-Realism emerges in this paper as a deliberate alternative to both techno-optimist inevitability and techno-pessimist withdrawal. Rather than assuming technological intervention as

157 inherently desirable or inherently corrupt, Techno-Realism treats intervention as morally contingent, institutionally
158 constrained, and politically situated. It asks not only whether technologies can be built, but whether they ought to be
159 built under specific social conditions, and who bears enduring responsibility for their consequences. This orientation
160 aligns with calls in responsible innovation and critical computing to foreground accountability, power, and long-term
161 governance rather than narrow notions of efficiency or usability [6, 34, 48].

163 Importantly, Techno-Realism does not reject technological design outright. Instead, it foregrounds conditionality,
164 recognizing that some technological interventions may be ethically justified when accompanied by robust governance,
165 community authority, and institutional accountability, while others warrant refusal, delay, or withdrawal. This stance
166 resonates with emerging scholarship that treats design as an ontological and political act shaping conditions of existence
167 rather than merely producing artifacts [24, 57]. Building on prior work that interrogates how problem–solution framings
168 in HCI reproduce racialized abstractions and ontological harm [18], Techno-Realism advances a normative framework
169 that explicitly evaluates the limits of technological intervention itself.

172 By situating itself between optimism and pessimism, Techno-Realism fills a conceptual gap in design scholarship: it
173 provides a structured mechanism for exercising restraint, accountability, and refusal without abandoning engagement
174 with real-world institutions, infrastructures, and policy regimes. In doing so, the framework enables designers and
175 researchers to navigate the ethical complexity of innovation under contemporary political–economic conditions, where
176 technological systems increasingly mediate governance, labor, and social life.

182 2.3 The Culture of Technology Innovation in HCI

183 Beyond political economy, techno-solutionism is also sustained through the everyday cultural practices, epistemic
184 norms, and evaluative standards of design communities. Understanding techno-solutionism in HCI therefore requires
185 situating innovation as a cultural practice, not a purely technical endeavor. Arnold Pacey’s work on the culture of
186 technology emphasizes that technological development reflects cultural values, social norms, and political priorities;
187 innovation cannot be disentangled from the contexts in which it is imagined, enacted, and evaluated [40, 41]. Pacey
188 argues that technology is expressive of cultural commitments and that engineering practices carry moral and political
189 implications, challenging the view that innovation is value-neutral or universally beneficial.

192 Within HCI, scholars have similarly articulated how design cultures embed specific assumptions about agency,
193 progress, and human needs. Stolterman highlights that design practice is shaped by epistemic frames that privilege
194 solvable problems and measurable outcomes, which in turn influences what counts as meaningful research and
195 innovation [52]. Bell and Dourish’s analysis of “My problem is your problem” illustrates how design communities
196 negotiate boundaries around what constitutes a legitimate problem—often privileging abstraction and generalizability
197 over situated, structural complexity [10]. DIS research has extended these critiques by showing how cultural valorizations
198 of disruption and novelty can eclipse deeper engagement with context, history, and equity in design processes [6, 29].

201 Taken together, these perspectives expose a culture of technology innovation that foregrounds intervention and
202 forward momentum while eliding the moral and historical dimensions of technological adoption. This elision contributes
203 to recurring patterns in which interventions reproduce the very conditions they intend to mitigate, particularly when
204 designers fail to contend with underlying power structures and institutional incentives that shape both problem
205 definitions and technological trajectories.

2.4 Human-Centered Computing and the Problem–Solution Binary

Human-computer interaction as a field is deeply rooted in the problem–solution binary, where research often begins by identifying a discrete problem and ends with a designed solution. This framing aligns with traditional engineering logics and positivist research paradigms, situating solvability and optimization as central goals of the discipline [39]. While this orientation has enabled advances in usability, accessibility, and user experience, it also constrains how social issues—especially structural injustices—are conceptualized and addressed.

Design scholars have critiqued this binary for overlooking the complexity of “wicked problems,” which lack definitive formulations, stable solutions, and uncontested success criteria [46]. Instead of converging on solutions, wicked problems often reveal deeper conflicts in values, power, and social organization. In HCI, this realization has led to work on anti-solutionism, which challenges the assumption that all design problems are tractable through technological intervention alone [12]. Complementary critiques highlight how the problem–solution frame can inadvertently naturalize social inequalities as design challenges to be corrected through individual behavior change or optimization, rather than as manifestations of structural dynamics requiring collective, political, and systemic responses [18, 56].

Extending this critique, recent theoretical work argues that the problem–solution pairing itself may be a “dead end” for research on equity and justice in computing. Drawing on philosophy and critical theory, this work demonstrates that the logic that binds problems to solutions can inadvertently perpetuate the conditions of harm it purports to address, particularly in contexts shaped by historical and ongoing inequalities [19]. From this perspective, confronting the limits of the problem–solution binary is not merely a methodological refinement but an epistemic shift toward frameworks that recognize the normative stakes of innovation and the potential necessity of refusal, withdrawal, or non-intervention.

2.5 Justice-Oriented Design and Computing Innovation Frameworks

HCI and adjacent fields have developed a robust set of justice-oriented design frameworks that contest extractive innovation and re-center those most affected by technology. Among the most influential, *Design Justice* articulates a community-led praxis for redistributing design power, challenging structural oppression, and building accountable design processes grounded in social movements and lived experience [16]. In contrast to universalist models of “the user,” *Design Justice* treats communities as epistemic authorities and insists that design outcomes and benefits should accrue to those historically marginalized by dominant design institutions.

Within CSCW and HCI, scholarship on community-based and participatory design has similarly emphasized relational accountability, long-term partnership, and the politics of participation. Harrington, Erete, and Piper show how community-based collaborative design is frequently constrained by institutional norms, timelines, and extractive research practices, and they call for more equitable participatory engagements that move beyond transactional models of “inclusion” [29]. Related work foregrounds intersectionality as a necessary analytic for designing with communities navigating multiple, overlapping systems of oppression, complicating simplistic notions of “marginalized users” and underscoring how design benefits and burdens are unevenly distributed [23]. These justice-oriented approaches have expanded what counts as legitimate design work in HCI: relationship-building, community governance, and collective sensemaking become central design activities rather than peripheral ethics.

Alongside these participatory traditions, values-based frameworks have sought to integrate normative commitments into design decisions. Value Sensitive Design (VSD) formalizes methods for surfacing stakeholder values and translating them into design requirements, offering a structured approach to ethical reflection in socio-technical systems [26, 27]. In

261 parallel, critical approaches to AI and socio-technical systems argue that fairness and accountability cannot be achieved
262 through technical metrics alone because abstraction often strips systems of their social meaning and institutional
263 embeddedness [48]. These lines of work provide essential foundations for justice-oriented innovation in computing.
264

265 However, even as these frameworks redistribute power and foreground values, many still presume that technology
266 remains an appropriate or promising site of intervention. In practice, justice-oriented design in HCI frequently retains
267 the design imperative—the assumption that the ethical response to injustice is to design *better* systems, rather than to
268 interrogate whether technological intervention is morally permissible, structurally capable, or politically legitimate
269 under conditions of capitalization and institutional constraint. This tension motivates approaches that treat refusal,
270 non-design, and withdrawal not as failures of innovation but as credible design outcomes when systems are likely to
271 compound harm.
272
273

274 2.6 Refusal, Prefiguration, and “Not-Doing” as Design Practice

275

276 A distinct and growing body of HCI and DIS scholarship challenges the presumption that design must culminate in new
277 technological artifacts. Rather than treating non-intervention as absence, this work conceptualizes “not-doing” as a
278 generative design stance: a way to contest the inevitability of technology and to render visible the political commitments
279 embedded in innovation. Pierce’s formulation of *undesigning* explicitly develops a framework for the intentional
280 negation of technology as a legitimate and productive aspect of design research, including inhibiting, displacing, or
281 removing technologies that produce negative effects [42]. In DIS, Sicart and Shklovski critique solutionist inclinations
282 by exposing how technological “solutions” can be ridiculous or misdirected when framed against imaginary or poorly
283 understood problems, underscoring the need to interrogate how problems are constituted through design itself [51].
284
285

286 In CSCW and HCI, prefigurative approaches provide another pathway for resisting dominant innovation logics.
287 Asad’s concept of *prefigurative design* argues that justice-oriented research should not merely study or mitigate inequity,
288 but should enact alternative relations and practices in the present, including community autonomy, accountability, and
289 transformation as methodological commitments [5]. This orientation aligns with broader calls in critical computing
290 to bring political and cultural critique into technical practice, rather than treating critique as external commentary
291 [2, 55]. Together, undesign, prefiguration, and critical technical practice expand the vocabulary of legitimate design
292 action: designers may build, but they may also refuse, withdraw, reallocate resources, redesign institutions, or change
293 governance arrangements.
294
295

296 These literatures are crucial for the Techno-Realist Design Framework because they legitimate two claims: (1) the
297 ethical response to harm may be *not* to build, and (2) justice-oriented innovation requires confronting the political
298 economy and institutional settings in which technologies are proposed and justified. Yet what remains under-specified
299 is a cohesive, practice-oriented framework for making these determinations in everyday design work, especially across
300 domains where designers are embedded in corporate, governmental, or academic systems with strong incentives to
301 deploy.
302
303

304 2.7 Gaps and Motivation: Why Techno-Realism is Needed

305

306 Taken together, justice-oriented design frameworks, participatory traditions, and refusal-oriented approaches provide
307 strong foundations for equity in HCI. Design Justice offers a movement-grounded praxis for redistributing power and
308 centering marginalized knowledge [16]; community-based collaborative design research clarifies the relational and
309 institutional conditions required for equitable participation [29]; values-based approaches formalize ethical deliberation
310 in design [26, 27]; and FAccT scholarship shows why technical abstraction repeatedly fails to capture the socio-technical
311

313 nature of harm [48]. Nonetheless, HCI still lacks a consolidated framework for addressing three recurring challenges in
314 justice-oriented innovation.
315

316 First, existing approaches often do not provide evaluative tests of moral permissibility for determining when techno-
317 logical intervention is justified, as opposed to merely preferable, inclusive, or well-intentioned. Second, accountability
318 is frequently framed at the level of participation or research practice, rather than as a lifecycle obligation spanning
319 design, deployment, governance, and remediation across institutions that hold power over technology trajectories [29].
320 Third, while refusal and undesign are increasingly recognized as legitimate stances, HCI lacks actionable guidance
321 for operationalizing refusal and withdrawal in settings where designers are constrained by capitalization, policy,
322 procurement, and institutional performance demands [42, 51].
323

324 The Techno-Realist Design Framework is motivated by these gaps. It complements Design Justice and participatory
325 traditions by foregrounding the limits of technological intervention under structural inequality and the political economy
326 of innovation, while offering actionable commitments for evaluating moral permissibility, allocating responsibility,
327 and operationalizing accountability beyond the artifact. In particular, Techno-Realism treats refusal, non-design, and
328 withdrawal as outcomes that must be supported by concrete organizational and policy mechanisms, rather than left as
329 rhetorical gestures. This framing provides a bridge between justice-oriented design commitments and the institutional
330 realities that routinely shape which technologies get built, for whom, and with what long-term consequences. The
331 framework operationalizes these commitments through a set of evaluative tests that generate concrete artifacts for
332 accountability, governance, and institutional decision-making.
333
334
335

336 **3 Methods: Auto-ethnographic Synthesis of Justice-Oriented Design Research**

337 This paper employs an auto-ethnographic method to synthesize over five years of justice-oriented design research
338 conducted across academic, industry, civic, and community settings. Rather than positioning auto-ethnography as
339 personal narrative, this work treats it as a rigorous, longitudinal, and reflexive mode of inquiry through which design
340 knowledge, normative commitments, and methodological insight emerge through practice. This approach enables
341 theoretical conceptualization grounded in sustained engagement with socio-technical systems and their consequences,
342 particularly where conventional empirical methods struggle to capture structural, ethical, and ontological dimensions
343 of design.
344
345
346

347 Auto-ethnography is used here not to foreground individual experience, but to trace how design decisions, institutional
348 constraints, community relationships, and accountability practices co-evolve over time. This orientation aligns with
349 prior work in HCI and CSCW that treats reflexivity, situated practice, and embodied judgment as legitimate sources of
350 design knowledge rather than methodological liabilities.
351

352 **3.1 Why Auto-ethnography in Design Research**

353 Auto-ethnography has an established lineage within HCI, CSCW, and design research as a method for examining
354 situated practice, reflexivity, and the researcher's entanglement within socio-technical systems [13, 20, 22]. In contrast
355 to detached observation, auto-ethnographic approaches foreground how knowledge is produced through participation,
356 negotiation, and institutional positioning, allowing for critical analysis of power, responsibility, and constraint.
357
358

359 This method is particularly appropriate for justice-oriented design research for three reasons. First, the work
360 synthesized here unfolds longitudinally across multiple projects, organizations, and community contexts, making short-
361 term studies insufficient for capturing how commitments evolve and how harms accumulate or shift over time. Second,
362 justice-oriented inquiry requires sustained reflexive engagement with values, accountability, and ethical tradeoffs that
363
364

are often invisible to conventional evaluative metrics. Third, the contribution of this paper is a design framework rather than a single artifact or intervention; such frameworks emerge through accumulated judgment and iterative practice rather than controlled experimentation alone.

Prior work at DIS, CHI, and CSCW increasingly recognizes auto-ethnographic and reflective methods as appropriate for theorizing design practice, particularly in domains involving care, ethics, and social justice [7, 25, 53]. Building on this tradition, the present study uses analytic auto-ethnography to surface recurring tensions, decision patterns, and normative commitments that structure justice-oriented design work in practice.

3.2 Data Sources

The analytic material for this study draws from multiple interrelated sources of design practice conducted between 2019 and 2024. These sources are treated as empirical traces of practice rather than autobiographical narrative.

Community Advisory Boards (CABs). Sustained collaboration with Black community-based organizations through the establishment and operation of a university-affiliated Community Advisory Board informed how accountability, power-sharing, and refusal can be operationalized within research institutions. These engagements included recurring meetings, co-developed agendas, governance agreements, and collective deliberation about technology impacts and research priorities [1].

Empirical Studies of Interactive AI Systems. Qualitative and mixed-method studies of digital financial platforms and automated speech recognition systems contributed insight into how techno-solutionist assumptions manifest in real-world deployment, particularly when systems framed as inclusive reproduce harm through data practices, evaluation protocols, or institutional incentives [1].

Design Workshops and Participatory Engagements. Justice-oriented design workshops, co-design sessions, and sustained dialogue processes provided experiential data on how communities articulate harm, value, and responsibility outside narrow problem–solution framings. These engagements informed the framework’s emphasis on relational accountability and ontological impact rather than solely functional outcomes [1].

Civic, Policy, and Institutional Engagement. Participation in public-interest technology initiatives, advisory activities, and policy-facing collaborations surfaced how design decisions intersect with regulation, procurement, organizational risk, and political economy. These contexts informed the framework’s attention to responsibility beyond individual designers and toward institutional governance structures [1].

Together, these sources capture a heterogeneous but coherent body of practice through which recurring ethical tensions, structural constraints, and accountability challenges became visible across domains.

3.3 Analytic Approach

Analysis proceeded through an iterative process of reflective synthesis rather than coding-driven thematic analysis. Drawing on analytic auto-ethnography [4] and reflective design research traditions [49], the analysis involved four interrelated analytic moves.

First, *reflective memos* documented moments of tension, failure, and ethical uncertainty across projects, particularly where conventional human-centered methods proved insufficient or misaligned with community priorities. Second, *cross-case patterning* identified recurring structural dynamics—such as misallocated responsibility, extractive participation, symbolic inclusion, and institutional inertia—that cut across domains and project types. Third, *normative abstraction* translated these patterns into explicit ethical commitments, clarifying what justice-oriented design demands beyond

417 usability, adoption, or fairness metrics. Finally, *framework construction* organized these commitments into a coherent
418 set of evaluative principles and procedural questions capable of guiding future design decision-making.

419 Critically, the Techno-Realist Design Framework did not precede the empirical work; it emerged through sustained
420 engagement with practical constraints, contradictions, and accountability failures encountered in situated design
421 contexts. This grounding in practice distinguishes the framework from speculative or purely theoretical approaches to
422 responsible design.
423

424 Taken together, this analytic process produced a set of recurring normative tensions and decision patterns that could
425 not be resolved through existing human-centered design heuristics alone. These tensions—concerning moral permis-
426 sibility, responsibility allocation, institutional constraint, community authority, and long-term accountability—were
427 abstracted into a structured framework intended to guide future design judgment rather than optimize technical
428 performance.
429

430 To make this analytic process transparent, traceable, and empirically grounded, the following subsection presents a
431 set of illustrative auto-ethnographic vignettes drawn from the underlying corpus of design practice. These vignettes
432 are not offered as representative cases or evidentiary proof in a positivist sense, but as analytic anchors that surface
433 recurring tensions, breakdowns, and ethical dilemmas encountered in justice-oriented design work. Each vignette
434 foregrounds a concrete moment of practice and traces how reflective synthesis translated situated experience into
435 normative commitments and evaluative concerns.
436

437 Together, the vignettes serve three purposes: (1) they demonstrate how abstract design commitments emerge from
438 lived institutional and community engagement rather than speculative reasoning; (2) they support traceability between
439 empirical experience and the framework articulated in Section 4; and (3) they invite readers to evaluate the interpretive
440 logic through which practice becomes theory. The vignettes should therefore be read as analytic exemplars rather than
441 illustrative anecdotes.
442
443
444

445 3.4 Illustrative Auto-ethnographic Vignettes

446 To ground the Techno-Realist Design Framework in empirical practice, this subsection presents three analytic vignettes
447 adapted from longitudinal auto-ethnographic field materials, including reflective memos and community-engaged
448 design work. The vignettes retain a reflexive voice consistent with auto-ethnographic traditions while anonymizing
449 institutional identifiers, locations, and project-specific details to preserve double-blind review. Each vignette follows
450 a structured analytic arc—context, tension, interpretation, and framework linkage—demonstrating how recurring
451 breakdowns and ethical dilemmas encountered in practice informed the formulation of specific design commitments
452 and evaluative tests.
453
454
455
456

457 *Vignette 1: Financial Access and the Limits of “Inclusion”.* *Context* In early doctoral work, I conducted a qualitative
458 study of Black American users’ experiences with a widely used peer-to-peer financial application supporting everyday
459 payments, remittances, and informal economic exchange [1]. Recruitment yielded several hundred prospective partici-
460 pants in a short period, signaling both high platform penetration and strong demand for articulating concerns about
461 digital financial infrastructures.
462

463 *Tension* Participants described account freezes, opaque fraud detection flags, and limited avenues for recourse when
464 funds were withheld. Although the platform positioned itself as expanding access and inclusion, many participants
465 experienced the system as fragile, punitive, and asymmetrically accountable. In design and policy discourse surrounding
466 the platform, these breakdowns were often framed as acceptable tradeoffs for fraud prevention and scalability. For
467
468

469 participants, however, system failures produced cascading harms: missed rent payments, disrupted family transfers,
470 and heightened mistrust toward financial institutions and digital mediation.

471 *Interpretation* This tension exposed a misalignment between optimization-driven design logics and lived financial
472 vulnerability. The core issue was not reducible to interface usability or model tuning, but whether the intervention was
473 morally defensible given the asymmetric distribution of risk and the weakness of remediation pathways. “Inclusion”
474 alone did not constitute an ethical justification for deployment when accountability mechanisms failed to protect those
475 most exposed to harm.
476

477 *Framework linkage* This vignette informed the commitments to *Normative Guidance* and *Recognizing Social Structures*.
478 It motivated the permissibility test as a mechanism for evaluating whether a system should be deployed under specific
479 social conditions and strengthened the requirement that designers analyze how technologies redistribute vulnerability,
480 institutional power, and exposure rather than treating inequity as a localized defect.
481
482

483 *Vignette 2: Community Advisory Work and the Burden of Participation. Context* As part of a university-affiliated
484 initiative, I established a community advisory board (CAB) to guide technology research priorities, ethical decision-
485 making, and participatory processes with community leaders and staff from community-based organizations. The
486 CAB convened over an extended engagement period and surfaced topic priorities that spanned surveillance, language
487 technologies, data governance, and other high-stakes systems.
488

489 *Tension* While participants expressed strong interest in shaping research agendas, sustaining meaningful partici-
490 pation required substantial coordination labor: compensation logistics, facilitation preparation, agenda-setting, and
491 ongoing relationship work. Community partners faced capacity constraints while balancing advisory participation with
492 organizational responsibilities. Academic timelines, funding structures, and administrative processes offered limited
493 support for this relational labor, creating uneven burdens across stakeholders and uncertainty about how community
494 authority would translate into durable decision rights.
495

496 *Interpretation* The experience clarified that participation alone does not guarantee equity or sustainability. Without
497 formal governance authority, durable resourcing commitments, and enforceable accountability mechanisms, partici-
498 patory initiatives can reproduce extractive dynamics even when justice-oriented intentions are present. Community
499 members may contribute significant labor while retaining limited leverage over research trajectories, dissemination,
500 and institutional decision-making.
501

502 *Framework linkage* This vignette shaped the commitments to *Localized Community Collaboration* and *Operationalizing*
503 *Accountability*. It motivated the governance threshold test, which evaluates whether communities possess meaningful
504 decision authority rather than symbolic inclusion, and strengthened the need for accountability artifacts (e.g., governance
505 agreements, compensation protocols, decision logs, and escalation pathways) that persist beyond interpersonal trust.
506
507
508
509

510 *Vignette 3: Dialogue-Based Design and Ontological Boundaries. Context* Before and during dissertation research, I
511 facilitated sustained dialogue workshops that brought participants together to engage issues of identity, power, and
512 conflict, with the aim of building shared understanding and enabling collective action. In technology-adjacent contexts,
513 I experimented with translating these practices into participatory design settings as a way to deepen interpretive
514 grounding before moving toward design proposals.
515

516 *Tension* Across workshops, participants frequently resisted premature translation of deeply personal or historically
517 grounded concerns into design requirements or feature proposals. Efforts to “design solutions” sometimes flattened
518 lived experience into abstractions that participants perceived as mis-representative or instrumentalizing. In multiple
519

521 instances, participants explicitly questioned whether technological intervention was appropriate at all, and whether
522 attempts to render identity or harm computationally legible risked reproducing the very forms of power being critiqued.

523 *Interpretation* These encounters highlighted ontological harm as distinct from functional failure. Even participatory
524 processes can reproduce harm when they render identity, trauma, or community knowledge legible primarily through
525 computational abstractions. Some social conditions resist technological mediation without loss of dignity, agency, or
526 epistemic integrity; in these cases, refusal or non-intervention may be the most ethical design outcome.

527
528 *Framework linkage* This vignette informed the commitment to *Addressing Ontological Harms* and reinforced refusal as
529 a legitimate design outcome. It motivated the ontological harm boundary test, which evaluates whether representational
530 or classificatory systems risk constituting harm regardless of performance metrics or participatory legitimacy.
531
532

533
534
535 *Summary.* Across these vignettes, recurring tensions—between optimization and vulnerability, participation and
536 power, and representation and dignity—exposed the limits of default solution-oriented design logics. The Techno-Realist
537 Design Framework consolidates these empirical insights into actionable commitments and evaluative tests that support
538 accountable decision-making, including when non-intervention, delay, or withdrawal is ethically warranted.
539

540 In the sections that follow, the Techno-Realist Design Framework is introduced as the formal articulation of these
541 analytic insights. The framework translates observed patterns from practice into explicit design commitments and
542 evaluative tests that can be operationalized within real-world design workflows, institutional/organizational review
543 processes, and governance contexts.
544
545

546 547 548 **4 The Techno-Realist Design Framework**

549 The Techno-Realist Design Framework offers a normative and practice-grounded approach to justice-oriented technology
550 research and design. It responds to the persistent limitations of techno-solutionist innovation by foregrounding realism
551 about what technology can, cannot, and should not do under conditions of structural inequality, institutional constraint,
552 and capitalist incentive. Rather than prescribing a universal design method or advocating for technological abstention,
553 Techno-Realism provides a set of design commitments—paired with evaluative tests—that guide designers in assessing
554 the moral permissibility, accountability, and social consequences of technological intervention.
555

556 Unlike reflective heuristics or ethical checklists, the evaluative tests in the Techno-Realist Design Framework
557 function as procedural checkpoints embedded within design and innovation workflows. Each test produces a concrete
558 artifact—such as a justification memo, responsibility map, governance agreement, or audit protocol—that enables
559 traceability, institutional memory, and accountability across the technology life-cycle. These tests can be applied during
560 project scoping, funding justification, deployment review, and post-deployment audit, allowing design decisions to be
561 documented, revisited, contested, and revised over time rather than treated as one-time ethical exercises.
562

563 The framework emerges from longitudinal engagement across community-based research, institutional collabora-
564 tions, and policy-facing design practice (Section 3). Across these contexts, recurring tensions surfaced: pressures to
565 deploy technologies despite unresolved harms; participation without decision-making power; accountability without
566 enforcement; and innovation framed as inevitable rather than contingent. Techno-Realism consolidates these tensions
567 into a coherent design stance that treats technological intervention as a choice requiring justification, rather than as an
568 assumed good.
569
570
571

Design Commitment	Evaluative Test	Artifact / Outcome
Normative Guidance	Is this technological intervention morally permissible in this context, or does it require refusal, delay, or remediation?	Written justification for intervention <i>or</i> documented decision to refuse or defer deployment
Assuming Responsibility	Who bears responsibility for harms across the technology lifecycle, and how is that responsibility enforced?	Responsibility map identifying accountable actors, obligations, and escalation paths
Recognizing Social Structures	How does this system interact with existing distributions of power, labor, risk, and harm?	Structural impact analysis detailing beneficiaries, exposed groups, and harm concentration
Localized Community Collaboration	Are affected communities engaged as co-governors with decision authority rather than symbolic participants?	Evidence of community decision-making power (e.g., veto rights, governance roles, agenda-setting authority)
Addressing Ontological Harms	How does the technology shape recognition, identity, and conditions of being beyond functional outcomes?	Identification of ontological or identity-based harms that may override performance or efficiency gains
Operationalizing Accountability	What concrete mechanisms enable oversight, contestation, repair, and withdrawal over time?	Defined accountability mechanisms (audits, withdrawal clauses, compensation pathways, oversight bodies)
Policy and Institutional Linkage	How are design decisions coupled to organizational policy, regulation, and institutional governance?	Explicit linkage to enforceable policies, regulatory requirements, or institutional governance structures

Table 1. Design commitments, evaluative tests, and resulting artifacts in the Techno-Realist Design Framework.

4.1 Framework Overview

At its core, the Techno-Realist Design Framework is defined by seven interrelated design commitments (Table 1). Each commitment articulates a normative orientation toward responsible innovation, while its associated test provides a practical lens for evaluating whether that commitment is being met in design practice. Together, these commitments shift design from a narrow problem-solving orientation toward a responsibility-centered practice that attends to history, power, and long-term consequences. Importantly, the framework recognizes refusal, non-design, delay, and withdrawal as legitimate—and sometimes necessary—design outcomes when technologies risk reproducing or intensifying harm.

In this framework, evaluative tests are not intended as reflective prompts alone, but as structured decision-making devices that guide action across the technology lifecycle. Each test functions as a justificatory threshold: it can be satisfied, partially met, or failed, with implications for whether a technological intervention proceeds, is revised, or is refused altogether. Tests require designers and institutions to produce and document forms of reasoning, evidence, and accountability—such as responsibility mappings, structural impact analyses, or governance commitments—that can be contested by affected stakeholders. In this sense, Techno-Realism treats evaluation not as measurement, but as situated judgment with consequences.

The framework is organized around seven design commitments derived from cross-case analytic synthesis. Each commitment is paired with an evaluative test that operationalizes how the commitment can be examined during project scoping, funding justification, deployment review, and post-deployment audit. These tests generate concrete artifacts and documentation that support traceability, accountability, and institutional memory across the life-cycle of a system.

Taken together, the commitments operate sequentially and recursively rather than as a static checklist. The permissibility test establishes whether technological intervention is justified at all; responsibility mapping and structural analysis locate accountability and systemic exposure; governance thresholds and ontological boundaries constrain acceptable

625 forms of impact; accountability mechanisms enable enforcement and repair; and policy linkage ensures durability
626 beyond individual projects or teams. In practice, these tests may be revisited iteratively as institutional conditions,
627 stakeholder relationships, and system impacts evolve, allowing the framework to function as a living governance
628 instrument rather than a one-time design intervention.
629
630

631 4.2 Design Commitments and Evaluative Tests 632

633 *Normative Guidance.* Techno-Realism begins from the commitment that not all technological interventions are
634 morally permissible. Rather than assuming innovation as inherently beneficial, designers are encouraged to evaluate
635 whether a proposed system meaningfully addresses harm or risks compounding it. This commitment draws on critiques
636 of solutionism and techno-optimism, which show how technologies framed as neutral or benevolent often stabilize
637 existing inequities and institutional power relations [14, 37].
638

639 The associated test asks designers to articulate the ethical justification for intervention and to consider refusal, delay,
640 or withdrawal as legitimate outcomes when harms cannot be adequately mitigated. Normative evaluation is ongoing
641 rather than front-loaded, recognizing that harms may only become visible once systems are embedded in lived contexts
642 [3]. Under Techno-Realism, the persistence of harm triggers obligations of remediation or withdrawal, not merely
643 iterative optimization.
644
645

646 *Assuming Responsibility.* Responsibility in Techno-Realism extends beyond individual designers to include institutions,
647 funders, vendors, and policymakers who shape technological trajectories. This commitment responds to documented
648 gaps in accountability where participatory or user-centered processes coexist with diffuse or unenforceable responsibility
649 [29].
650

651 The evaluative test requires designers to make responsibility explicit: who is accountable for harm, how responsibility
652 is distributed across the lifecycle of a system, and what mechanisms exist for redress when systems fail. Responsibility
653 persists beyond deployment, particularly when harms disproportionately affect historically marginalized communities
654 [8]. Transparency, grievance processes, and enforceable remediation pathways are treated as design requirements rather
655 than optional governance add-ons.
656
657

658 *Recognizing Social Structures.* Technological systems are embedded within social, economic, and political structures
659 that shape how benefits and harms are distributed. This commitment requires designers to analyze how systems interact
660 with histories of exclusion, labor exploitation, and governance, rather than treating harm as an isolated bias or technical
661 flaw [48, 58].
662

663 The associated test asks whether a technology reorganizes power, risk, or access in ways that exacerbate structural
664 inequality. Crucially, Techno-Realism emphasizes that the intended user is often not the most affected subject of a
665 system's operation [15]. When structural harms exceed a system's capacity for ethical governance, refusal or moratoria
666 become appropriate design responses.
667
668

669 *Localized Community Collaboration.* While justice-oriented design emphasizes participation, Techno-Realism dis-
670 tinguishes between inclusion and governance. This commitment asserts that communities should not merely provide
671 feedback but participate as collaborators with authority to shape research agendas, evaluation criteria, and decisions to
672 halt or redirect design work [16, 25].
673

674 The evaluative test asks whether participatory mechanisms confer real decision-making power or function primarily
675 as symbolic inclusion. This distinction responds to critiques of participatory design that show how engagement
676

677 can coexist with extractive or predetermined innovation trajectories [28, 30]. Under Techno-Realism, community
678 collaboration is meaningful only when it can materially influence outcomes.
679

680
681 *Addressing Ontological Harms.* Beyond functional performance, technologies shape who is recognized, categorized,
682 and rendered legible. Techno-Realism foregrounds ontological harm: harm to identity, dignity, and conditions of being
683 that cannot be captured through accuracy, efficiency, or fairness metrics alone [15, 38, 56].
684

685 The evaluative test asks how a system constructs normalcy, deviance, legitimacy, or risk, and whether these con-
686 structions foreclose alternative ways of being. Ontological harm may persist even in participatory or well-intentioned
687 systems; accordingly, its presence constitutes grounds for refusal regardless of procedural legitimacy or technical
688 performance.
689

690
691 *Operationalizing Accountability.* Accountability must be actionable rather than aspirational. This commitment
692 emphasizes the need for concrete mechanisms—such as community oversight bodies, auditability, withdrawal clauses,
693 and material compensation—that persist beyond initial deployment [34, 42].
694

695 The evaluative test examines whether designers and institutions have planned for failure, including pathways
696 for repair, redress, and decommissioning when systems cause irreparable harm. By operationalizing accountability,
697 Techno-Realism treats repair and withdrawal as ongoing responsibilities rather than exceptional responses.
698

699
700 *Policy and Institutional Linkage.* Finally, Techno-Realism rejects the notion that ethical design can be isolated at the
701 artifact level. This commitment insists that design decisions be explicitly coupled to organizational policy, procurement
702 practices, regulation, and governance structures [6].
703

704 The evaluative test asks whether institutional conditions exist to enforce ethical commitments over time. Many
705 harms attributed to technology are enabled not by design alone, but by policy vacuums or misaligned incentives [44].
706 Under Techno-Realism, engaging policy processes is an integral part of design responsibility.
707

708 709 710 **4.3 From Framework to Practice**

711 Techno-Realism does not prescribe a singular method or outcome. Instead, it offers a scaffold for deliberation, refusal,
712 and accountability that can be adapted across domains of interactive computing. By framing justice-oriented design as
713 a set of commitments subject to ongoing evaluation, the framework equips researchers and practitioners to decide not
714 only how to build technologies, but whether they should be built at all. In the following section, we discuss how this
715 framework extends prior justice-oriented approaches and consider its implications for HCI research and practice.
716

717 While Techno-Realism is informed by justice-oriented design traditions, it departs from them by centering moral
718 permissibility, refusal, and post-deployment obligation as first-class design concerns. Rather than assuming technological
719 intervention as an appropriate response once participation is achieved, the framework interrogates whether intervention
720 itself is justified and who bears responsibility when harm persists.
721

722 Techno-Realism does not resolve disagreement; rather, it makes disagreement visible and consequential. Evaluative
723 tests may surface irreconcilable positions among designers, institutions, and communities, particularly where harms
724 are unevenly distributed or politically contested. In such cases, the framework does not promise resolution but insists
725 on accountability for how decisions are made and justified.
726

4.4 Applied Vignettes: Applying the Normative Guidance (Permissibility) Test

Normative Guidance asks whether a technological intervention is morally permissible in a given context, or whether refusal, delay, or remediation is required. Consider a research–industry collaboration developing an AI-driven risk assessment tool intended to support hospital triage decisions. Early prototypes demonstrate high predictive accuracy on retrospective data, and participatory workshops with clinicians suggest the system could streamline decision-making under resource constraints. However, applying the normative guidance test requires the design team to evaluate permissibility beyond performance and participation. A structural analysis reveals that the training data encode racialized patterns of access to care, producing systematically higher risk scores for patients from historically underserved communities. Further engagement with patient advocates surfaces concerns that the tool would legitimize existing disparities by rendering them clinically “objective.”

At this stage, the evaluative test is not satisfied. The design team documents a formal justification for delaying deployment, citing the system’s role in stabilizing inequitable care pathways rather than mitigating them. Despite institutional pressure to proceed—given funding timelines and apparent technical success—the test triggers a decision to pause development and pursue remediation. This includes revisiting the framing of “risk,” exploring non-algorithmic alternatives, and negotiating with institutional partners about conditions under which the tool could be ethically reconsidered. In this vignette, the normative guidance test functions as a justificatory threshold: it produces a documented decision to delay intervention, foregrounds responsibility for harm, and legitimizes refusal as a valid design outcome rather than a failure of innovation.

5 Discussion

This discussion situates the Techno-Realist Design Framework within existing HCI and design scholarship, clarifies its distinctive contributions, and reflects on its implications for research, practice, and technology development beyond academic settings. I argue that Techno-Realism extends prior justice-oriented and participatory frameworks by foregrounding evaluative judgment, refusal, and post-deployment responsibility as first-class design concerns. In doing so, it invites a rethinking of how problems, solutions, and innovation itself are framed within HCI and related fields.

5.1 How Techno-Realism Extends Prior Design Frameworks

The Techno-Realist Design Framework is informed by, but not reducible to, several established traditions in HCI and design research. Where prior frameworks emphasize participation, values articulation, or speculative critique, Techno-Realism foregrounds *evaluative decision-making under constraint*, particularly in institutional and high-stakes contexts.

Design Justice and related justice-oriented frameworks emphasize redistributing design power, centering marginalized knowledge, and ensuring meaningful participation in design processes [16]. Techno-Realism aligns with these commitments but extends them by explicitly interrogating the moral permissibility of technological intervention itself. Where Design Justice asks *who decides* and *who benefits*, Techno-Realism additionally asks *whether this technology should exist at all*, and *who remains responsible when harm persists despite participatory intent*. This shift is particularly salient in domains where inclusion alone cannot mitigate structural or ontological harm.

Participatory Design (PD) has long emphasized co-design, mutual learning, and democratic engagement [47]. Techno-Realism builds on PD’s commitment to shared authority while responding to critiques that participatory processes can coexist with predetermined outcomes, extractive relationships, or institutional inertia [28, 30]. By introducing

781 governance thresholds and responsibility chains, Techno-Realism re-frames participation as a question of decision
782 authority and enforceability rather than presence alone.

783 Value-Sensitive Design (VSD) foregrounds the systematic incorporation of human values throughout the design
784 lifecycle [27]. While VSD provides robust conceptual and methodological tools for surfacing values, Techno-Realism
785 diverges by treating evaluation not as alignment with articulated values, but as a situated judgment about consequences,
786 obligations, and failure conditions. In this sense, Techno-Realism is less concerned with value elicitation and more
787 concerned with moral thresholds and accountability when values come into conflict.

788
789 Finally, Critical and Speculative Design practices challenge dominant assumptions by imagining alternative futures
790 and exposing hidden politics of technology [21]. Techno-Realism complements these approaches by operating in
791 the register of intervention rather than provocation. Rather than speculating on alternatives, it asks how designers
792 act responsibly when intervention is already underway, funded, or institutionally mandated—and when refusal or
793 withdrawal may be the most ethical outcome.
794
795

796 797 **5.2 Rethinking Solutions, Problems, and Refusal**

798
799 A central contribution of Techno-Realism is its challenge to the problem–solution framing that dominates technology
800 design and innovation. HCI scholars have long critiqued techno-solutionism—the tendency to frame complex social
801 issues as solvable through technical artifacts [12, 18, 37]. Yet even critical approaches often retain the assumption that
802 problems must be addressed through design interventions.
803

804 Techno-Realism pushes this critique further by making space for refusal, delay, and non-intervention as legitimate
805 design outcomes. Drawing on ontological design scholarship, which emphasizes that design shapes conditions of being
806 rather than merely solving problems [24, 57], the framework treats technological solutions as world-making acts that
807 may foreclose alternative forms of existence. In this framing, harm is not only a matter of bias or misalignment, but of
808 how technologies define normalcy, risk, and legitimacy [15, 18, 56].
809

810 The framework’s evaluative tests operationalize this stance by allowing solutions to be examined without presuming
811 a corresponding problem. This move resonates with scholarship on design refusal and undesign, which argues that
812 restraint and withdrawal can be generative design acts rather than failures [9, 42]. Techno-Realism contributes to this
813 literature by embedding refusal within a broader evaluative doctrine—one that treats refusal as a consequence of failing
814 moral, structural, or ontological tests, rather than as an individual ethical stance [17].
815
816
817

818 **5.3 Implications for HCI Researchers and Designers**

819 For stewards of interactive computing systems, whether developers, researchers, or designers, Techno-Realism has
820 several practical implications for how projects are scoped, evaluated, and concluded.

821 First, the framework encourages researchers to treat project scoping as an ethical act. Rather than beginning with a
822 problem statement and assumed solution, researchers can use evaluative tests to interrogate whether technological
823 intervention is appropriate, and under what conditions it should proceed. This has implications for grant proposals,
824 industry partnerships, and student projects alike.
825

826 Second, Techno-Realism re-frames evaluation and success. Success is not defined solely by usability, adoption, or
827 technical performance, but by whether commitments around responsibility, structural impact, and accountability are
828 satisfied. In some cases, success may take the form of delaying deployment, narrowing system scope, or documenting
829 refusal.
830

833 Third, the framework pushes designers to involve communities beyond participation, toward shared governance and
834 decision authority. This requires rethinking timelines, compensation structures, and institutional incentives that often
835 marginalize community knowledge once design moves toward deployment.
836

837
838 *Implications for Software Ventures and Entrepreneurship.* These implications extend beyond academic HCI into the
839 landscape of software ventures and technology entrepreneurship. Venture-backed innovation environments are often
840 driven by imperatives to scale quickly, capture markets, and demonstrate growth—logics encapsulated in mottos such
841 as “move fast and break things.” As Ruha Benjamin asks in her critique of technological innovation, “what about the
842 people and places broken in the process?” [11].
843

844 The recent proliferation of generative AI applications illustrates this tension. Since 2022, a flood of under-specified,
845 low-value, and ethically questionable AI products have entered workplaces, schools, and public institutions, often
846 without clear demand or accountability. Scholars in information systems and technology management have noted that
847 such dynamics are reinforced by venture capital incentives that prioritize speed and novelty over long-term social
848 impact [35, 59].
849

850 Applied early in the lifecycle of software ventures, Techno-Realism offers a counterweight to these pressures.
851 Evaluative tests can function as stage-gates in product development, requiring founders, investors, and designers to
852 justify intervention, map responsibility, and anticipate harm before scaling. Rather than treating ethics as post hoc risk
853 management, Techno-Realism integrates moral permissibility and accountability into decisions about what products
854 are built, funded, and deployed in the first place.
855

856
857 *In summary.* These discussions position Techno-Realism as a framework for responsible decision-making in contexts
858 where design is entangled with institutional power, economic incentive, and social harm. By foregrounding evaluative
859 judgment, refusal, and accountability, the framework extends justice-oriented design traditions while addressing
860 persistent gaps in how technology design is governed in practice. For DIS and the broader HCI community, Techno-
861 Realism offers a way to move beyond better solutions toward more responsible decisions about when, how, and whether
862 to intervene at all.
863
864

865 5.4 Limitations and Future Work 866

867 This work advances Techno-Realism as a normative and practice-grounded framework rather than a universally
868 generalizable method or prescriptive design protocol. While the paper demonstrates how evaluative tests can be
869 operationalized through concrete artifacts and applied reasoning, the framework intentionally does not provide step-
870 by-step procedures, optimization criteria, or guarantees of consensus in contested design contexts. This is a deliberate
871 design choice: Techno-Realism foregrounds judgment, moral deliberation, institutional responsibility, and the legitimacy
872 of refusal rather than procedural compliance or technical optimization. As a result, the framework should be understood
873 as a scaffold for accountable decision-making rather than a deterministic instrument.
874

875
876 Empirically, the framework is derived from longitudinal auto-ethnographic synthesis across justice-oriented research
877 and institutional collaborations primarily situated within U.S.-based academic, nonprofit, and public-sector environments.
878 These engagements span multiple application domains, including language technologies, financial technologies, civic
879 data systems, and generative AI initiatives. However, the political, regulatory, cultural, and economic conditions shaping
880 these contexts are not universal. The framework may surface different tensions when applied in settings such as
881 consumer technology startups, Global South innovation ecosystems, safety-critical engineering, informal technology
882 practices, or jurisdictions with distinct governance regimes. Accordingly, Techno-Realism is offered for analytic transfer
883
884

rather than direct generalization, inviting empirical examination of how its commitments and evaluative tests operate under divergent institutional constraints and cultural assumptions.

Finally, while this paper focuses on theorizing evaluative commitments and their role in structuring accountable design judgment, additional work is needed to study how Techno-Realism is enacted in everyday organizational practice over time. Future research may explore how the framework can be translated into facilitation practices, documentation workflows, training processes, and governance mechanisms that support sustained institutional uptake without reducing the framework to compliance checklists or managerial instrumentation. Such work would enable systematic study of disagreement, breakdown, adaptation, and refusal in real-world settings, strengthening understanding of how justice-oriented design commitments are negotiated under practical constraints.

6 Conclusion

This paper introduced the Techno-Realist Design Framework as a normative and practice-grounded approach for evaluating technological intervention under conditions of structural inequality, institutional constraint, and capitalist incentive. Rather than offering another prescriptive design method, the framework re-frames justice-oriented innovation as a set of commitments and evaluative tests that guide when, how, and whether technologies ought to be built at all. By foregrounding moral permissibility, responsibility, structural analysis, and the legitimacy of refusal, Techno-Realism responds to long-standing critiques of techno-solutionism while remaining grounded in the realities of contemporary design practice.

The framework emerges from longitudinal, situated engagement across community-based research, institutional collaborations, and policy-facing design work, and is intended to travel across domains of interactive computing. In articulating evaluative tests alongside concrete artifacts and outcomes, Techno-Realism offers researchers and practitioners a scaffold for making accountable design decisions without assuming technological intervention as an inherent good. Importantly, it positions non-design, delay, and withdrawal as ethical outcomes rather than failures of innovation.

By situating design within broader systems of power, governance, and political economy, this work extends justice-oriented design discourse beyond participation and inclusion toward responsibility and constraint. Techno-Realism invites the DIS community to reconsider not only how technologies are designed, but how design decisions are justified, governed, and lived with over time. In doing so, it contributes a pragmatic yet critical framework for advancing human-centered computing that is accountable to the people and places most affected by technological change.

References

- [1] [n. d.]. Citation anonymized for double-blind review..
- [2] Philip E. Agre. 1997. Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI. In *Social Science, Technical Systems, and Cooperative Work: Beyond the Great Divide*, Geoffrey C. Bowker, Les Gasser, Susan Leigh Star, and William Turner (Eds.). Lawrence Erlbaum Associates, 131–157.
- [3] Morgan G. Ames. 2019. *The Charisma Machine: The Life, Death, and Legacy of One Laptop per Child*. MIT Press.
- [4] Leon Anderson. 2006. Analytic Autoethnography. *Journal of Contemporary Ethnography* 35, 4 (2006), 373–395.
- [5] Mariam Asad. 2019. Prefigurative Design as a Method for Research Justice. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3359302
- [6] Seyram Avle and Silvia Lindtner. 2018. Responsibility, Accountability, and Innovation in Human-Centered Computing. *ACM Transactions on Computer-Human Interaction* 25, 6 (2018), 1–36.
- [7] Shaowen Bardzell and Jeffrey Bardzell. 2018. Feminist HCI: Taking Stock and Outlining an Agenda for Design. *Human-Computer Interaction* 33, 4 (2018), 309–359.
- [8] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. Fairness and Machine Learning. *fairmlbook.org* (2019). Online manuscript.

- [9] Eric P. S. Baumer, M. Six Silberman, Sarah Fox, and Bonnie A. Nardi. 2014. Reflective Informatics: Conceptualizing Reflection in Information Systems Research. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)* (2014), 585–594.
- [10] Genevieve Bell and Paul Dourish. 2008. My Problem Is Your Problem: Negotiating Boundaries in CSCW. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*. ACM.
- [11] Ruha Benjamin. 2019. Race after technology: Abolitionist tools for the new jim code. *Social forces* (2019).
- [12] Mark Blythe, Enrique Encinas, and Andrew Monk. 2016. Anti-Solutionism. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM.
- [13] Kirsten Boehner, Janet Vertesi, Phoebe Sengers, and Paul Dourish. 2007. How HCI Interprets the Probes. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2007).
- [14] Meredith Broussard. 2019. *Artificial Unintelligence: How Computers Misunderstand the World*. MIT Press.
- [15] Simone Browne. 2015. *Dark Matters: On the Surveillance of Blackness*. Duke University Press, Durham, NC.
- [16] Sasha Costanza-Chock. 2020. Design Justice: Community-Led Practices to Build the Worlds We Need. *MIT Press* (2020).
- [17] Jay L. Cunningham. 2024. *Methods of Designing Justice-Oriented Interactive AI Systems*. Ph. D. Dissertation. University of Washington, Seattle, WA, USA. <https://digital.lib.washington.edu/server/api/core/bitstreams/f0ec376-55a4-4a82-9cb5-70f081e00e51/content>
- [18] Jay L. Cunningham, Gabrielle Benabdallah, Daniela K. Rosner, and Alex Taylor. 2023. On the Grounds of Solutionism: Ontologies of Blackness and HCI. *ACM Transactions on Computer-Human Interaction* 30, 2 (2023).
- [19] Jay L. Cunningham, Gabrielle Benabdallah, Daniela K. Rosner, and Alex Taylor. 2023. On the Grounds of Solutionism: Ontologies of Blackness and HCI. *ACM Transactions on Computer-Human Interaction* 30, 2 (2023).
- [20] Paul Dourish. 2014. Reading and Interpreting Ethnography. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2014).
- [21] Anthony Dunne and Fiona Raby. 2013. *Speculative Everything: Design, Fiction, and Social Dreaming*. MIT Press.
- [22] Carolyn Ellis, Tony E. Adams, and Arthur P. Bochner. 2011. *Autoethnography: An Overview*. Vol. 12. Forum: Qualitative Social Research.
- [23] Sheena Erete, Aarti Israni, and Tawanna Dillahunt. 2018. An Intersectional Approach to Designing in the Margins. *interactions* 25, 3 (2018), 66–69. doi:10.1145/3194349
- [24] Arturo Escobar. 2018. *Designs for the Pluriverse: Radical Interdependence, Autonomy, and the Making of Worlds*. Duke University Press.
- [25] Sarah Fox and Christopher A. Le Dantec. 2017. Community Historians: Scaffolding Community Engagement through Culture and Heritage. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW, Article 24 (2017).
- [26] Batya Friedman and David G. Hendry. 2019. *Value Sensitive Design: Shaping Technology with Moral Imagination*. The MIT Press, Cambridge, MA.
- [27] Batya Friedman, Peter H. Kahn, and Alan Borning. 2006. Value Sensitive Design and Information Systems. *Human-Computer Interaction* 21, 2 (2006), 241–273.
- [28] Kishonna L. Gray, Janet Vertesi, Alex Taylor, and Christina Harrington. 2020. Critical Race Theory for HCI. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020).
- [29] Christina N. Harrington, Sheena Erete, and Anne Marie Piper. 2019. Deconstructing Community-Based Collaborative Design: Towards More Equitable Participatory Design Engagements. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW, Article 216 (2019). doi:10.1145/3359318
- [30] Gillian R. Hayes. 2011. The Relationship of Action Research to Human-Computer Interaction. *ACM Transactions on Computer-Human Interaction* 18, 3 (2011).
- [31] Anna Lauren Hoffmann. 2021. Where Fairness Fails: Data, Algorithms, and the Limits of Antidiscrimination Discourse. *Information, Communication & Society* 24, 7 (2021), 900–915.
- [32] Michael H. Huesemann and Joyce A. Huesemann. 2011. *Techno-Fix: Why Technology Won't Save Us or the Environment*. New Society Publishers.
- [33] Lilly Irani. 2019. Chasing Innovation: Making Entrepreneurial Citizens in Modern India. *Princeton University Press* (2019).
- [34] Jade E. Irvine and Lilly Irani. 2019. Accountable Algorithms. *ACM Interactions* 26, 4 (2019).
- [35] Martin Kenney and John Zysman. 2014. *Silicon Valley Revisited: The New Entrepreneurial Region*. Stanford University Press.
- [36] Charlton McIlwain. 2019. *Black Software: The Internet & Racial Justice*. Oxford University Press.
- [37] Evgeny Morozov. 2013. *To Save Everything, Click Here: The Folly of Technological Solutionism*. PublicAffairs, New York, NY.
- [38] Safiya Umoja Noble. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press, New York, NY.
- [39] Antti Oulasvirta and Kasper Hornbæk. 2016. HCI as Optimization. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016).
- [40] Arnold Pacey. 1983. *The Culture of Technology*. MIT Press.
- [41] Arnold Pacey. 2001. *Meaning in Technology*. MIT Press.
- [42] James Pierce. 2012. Undesigning Interaction. In *Proceedings of the 2012 CHI Conference on Human Factors in Computing Systems*. ACM.
- [43] James Pierce. 2015. Working by not quite working. In *Proceedings of CHI*.
- [44] Iyad Rahwan, Manuel Cebrian, Nick Obradovich, Josh Bongard, Jean-Francois Bonnefon, Cynthia Breazeal, Jacob Crandall, Nicholas Christakis, Jeroen Van Den Hoven, Igor Fasel, et al. 2019. Machine Behaviour. *Nature* 568, 7753 (2019), 477–486.
- [45] Yolanda Rankin and Jakita Thomas. 2020. Intersectionality and HCI: Engaging race, gender, and class. In *Proceedings of CHI*.
- [46] Horst W. J. Rittel and Melvin M. Webber. 1973. Dilemmas in a General Theory of Planning. *Policy Sciences* 4, 2 (1973), 155–169.
- [47] Douglas Schuler and Aki Namioka. 1993. Participatory Design: Principles and Practices. In *Participatory Design: Principles and Practices*. Lawrence Erlbaum Associates.

- 989 [48] Andrew D. Selbst, Danah Boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and Abstraction in Sociotechnical
990 Systems. *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)* (2019).
- 991 [49] Phoebe Sengers, Kirsten Boehner, Shay David, and Joseph Kaye. 2005. Reflective Design. In *Proceedings of the SIGCHI Conference on Human Factors
992 in Computing Systems (CHI '05)*. ACM, New York, NY, USA.
- 993 [50] Miguel Sicart. 2020. Pataphysical design. *Proceedings of DIS* (2020).
- 994 [51] Miguel Sicart and Irina Shklovski. 2020. Pataphysical HCI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM,
995 Honolulu, HI, USA. doi:10.1145/3313831.3376760
- 996 [52] Erik Stolterman. 2008. The Nature of Design Practice and Implications for Interaction Design Research. *International Journal of Design* 2, 1 (2008).
- 997 [53] Michael Carl Tschantz and Batya Friedman. 2015. Designing Values. In *Proceedings of the 2015 CHI Conference on Human Factors in Computing
998 Systems*. ACM.
- 999 [54] Fred Turner. 2019. *The Democratic Surround: Multimedia and American Liberalism from World War II to the Psychedelic Sixties*. University of Chicago
Press.
- 1000 [55] Daniela van Geenen. 2024. Pluralising critical technical practice. *Subjectivity* 17 (2024), 1–20.
- 1001 [56] Calvin L. Warren. 2018. *Ontological Terror: Blackness, Nihilism, and Emancipation*. Duke University Press, Durham, NC.
- 1002 [57] Anne-Marie Willis. 2006. Ontological Designing. *Design Philosophy Papers* 4, 2 (2006), 69–92.
- 1003 [58] Langdon Winner. 1980. Do Artifacts Have Politics? *Daedalus* 109, 1 (1980), 121–136.
- 1004 [59] Eric Yanfei Zhao and Michael Lounsbury. 2015. Institutional Logics and the Dynamics of Innovation: A Perspective from Organization Studies.
1005 *Organization Studies* 36, 3 (2015), 363–390.

1006 Received 19 January 2026; revised 12 March 2009; accepted 5 June 2009

1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040

Manuscript submitted to ACM