*Cambridge Forum on AI: Culture and Society* (2025), **1**, e2, 1–8
doi:10.1017/cfc.2025.10004

Cambridge Forum

REFLECTION

# AI interdisciplinarity as critical technical practice

Lucy Suchman[1], Sireesh Gururaja[2] and David Gray Widder[3]

[1]Department of Sociology, Lancaster University, Lancaster, Lancashire, UK; [2]Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA and [3]Digital Life Initiative, Cornell Tech, New York, NY, USA
**Corresponding author:** Lucy Suchman; Email: l.suchman@lancaster.ac.uk

(Received 1 April 2025; revised 30 June 2025; accepted 15 July 2025)

**Abstract**

This Reflection draws from an ongoing collaboration among the three authors, investigating mutual enlistment between the United States Department of Defense and research communities devoted to the advancement of the AI project. We are interested in the intersecting concerns and resonant sensibilities that draw us together – what we argue is the necessary starting point for interdisciplinary thinking – and in the differences that are the collaboration's generative possibilities. Among the threads that join us are intersecting pathways between academic and commercial positionings within and against the AI project. Each of us has moved between locations in industry and the university, as researchers, practitioners, and students/faculty. Drawing on these experiences, we explore the possibilities that different positionings afford and what they preclude, how we have attempted to navigate these institutions within their frames of reference, and what has drawn us into relations beyond their putative boundaries. Based on Philip Agre's call for a "critical technical practice" as a path towards more radical shifts in knowledge practices, we consider how we might weave together our biographical trajectories, disciplinary affiliations, political commitments, subjectivities, and skills into what Andrew Barry and Georgina Born name a more "auspicious interdisciplinarity".

**Keywords:** AI critque; computing; interdisciplinarity; knowledge practices

The inaugural issue of *Cambridge Forum on AI: Culture and Society* calls on us to consider interdisciplinarity not as a general construct, but through relevant social histories and in specific sites of research practice. Resources for that programme include a rich and growing body of investigations into the entangled intellectual, technical, political, and economic lineages of the project of creating machines in the image of the human, abbreviated under the acronym of Artificial Intelligence (AI).[1] In the present moment, increasing numbers of us are engaged – whether collaboratively or agonistically – with that project.

This essay draws from an ongoing collaboration among the three authors. We reflect on the intersecting concerns and resonant sensibilities that brought us together and on the differences that are the collaboration's generative possibilities. We share some starting premises; for example, that technology is always already social relations, that "innovation" is less an inherent and distinguishing attribute than a product of differential valuation, and that accountability is a matter not of audit but of justice. Within our collaboration, AI is variously a technoscientific domain for critical inquiry, a

[1]For indicative histories, see Adam, 1998; Jones, 2016; Mackenzie, 2017; Pasquinelli, 2023; Riskin, 2007.

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/cfc.2025.10004 Published online by Cambridge University Press

field of professional and academic affiliation, and a site of technical practice. Disciplinary identifications (anthropology/science and technology studies (STS), computer science/software engineering) are clearly relevant here, including those that we take to be our frames of reference, those that define the communities of discourse and practice that we inhabit, and those to whom our work is addressed.

Among the threads that join us are intersecting travels between academic and commercial positionings, working both within and against the AI project. Each of us has moved between locations in industry and the university as researchers, practitioners, students and faculty. Drawing on these experiences, we explore the possibilities that those different locations afford and what they preclude, how we have attempted to navigate these institutions within their frames of reference, and what has drawn us into relations beyond their putative boundaries.

Our shared project is to explore modes of resistance to interdisciplinary practice that are subservient to dominant political and economic interests. In this sense, we align with the proposition that "ideas of interdisciplinarity imply a variety of boundary transgressions, in which the disciplinary and disciplining rules, trainings, and subjectivities given by existing knowledge corpuses are put aside" (Barry and Born, 2013:1). Our aim is not so much to put aside received knowledge as to explore the possibilities for generative, agonistic engagement with the multiplicity that constitutes a discipline—a form of what Philip Agre (1997a, b) names a "critical technical practice." With this aim in mind, we consider how we might weave together our biographical trajectories, disciplinary affiliations, political commitments, trainings, and subjectivities into a more "auspicious interdisciplinarity" (Barry and Born, 2013: 39).

## 1. The project that brought us together

As is the case for most research collaborations, the project that brought us together was formed not as an exercise in interdisciplinarity *per se*, but as an opportunity to deepen our respective engagements with a problem of shared interest. Our work together came about at a time characterized by the massive promotion of AI, not least in its application to the advancement of U.S. military supremacy. In the fall of 2023, David had just completed a PhD in Software Engineering at Carnegie Mellon University (CMU), which derives 42% of its research funding from the U.S. Department of Defense (DoD).[2] He had drawn on Lucy's work in an article building upon his PhD dissertation (Widder and Nafus 2023) to argue that the supply chain nature of AI systems contributes to the systematic abdication of responsibility by individual practitioners. David's subsequent writing (Widder, West, and Whittaker 2024), problematizing the trope of "open AI" and challenging claims that open-source software can "democratize" access to AI in the absence of more fundamental transformations, was another point of connection. These intersecting interests in AI critique comprised the starting place for our collaboration.

In 2023, David had obtained a substantial dataset comprising DoD grant solicitations issued from 2007 to 2023.[3] The question was whether and how this large document corpus might provide material relevant to an investigation of research and development in algorithmically enabled warfare, shedding further light on the relations that comprise the U.S. military-academic complex. Thinking about methods for engaging with the corpus, David contacted his friend Sireesh, then a third-year PhD student at CMU studying Natural Language Processing (NLP), whom he had met through graduate unionization efforts. Sireesh had previously taken a leading role in a qualitative study interviewing AI/NLP researchers on the incentives that shape their practice, aimed at generating a wider discussion of the field's implicit norms in the service of more deliberately (re)directing its

[2]Carnegie Mellon University. 2023, March 28. State of the University Presentation. https://www.cmu.edu/leadership/assets/pdf/state-of-the-university-presentation-032823.pdf.

[3]The original corpus comprised 46,175 solicitations automatically downloaded in bulk from the publicly available U.S. government grants.gov database, which we filtered for those originating from the Department of Defense, narrowing our corpus down to create a final dataset of 7,187 documents.

https://doi.org/10.1017/cfc.2025.10004 Published online by Cambridge University Press

future (Gururaja et al 2023). Sireesh was also, at the time, funded by a DoD grant himself and eager to deepen his understanding of the implications of that position. As well as becoming a co-author, Sireesh designed a customized search interface for the dataset tailored to our interests, enabling us to focus on documents within the corpus relevant to research in AI and Machine Learning. The project, in this sense, was at once a critical analysis of AI funding and an exercise in the development of NLP-enabled search to support that analysis.

The analysis that resulted was framed by a critique of the distinction between basic and applied research, showing how funding calls designated as basic research nonetheless enroll researchers into a warfighting agenda (Widder, Gururaja, and Suchman 2024).[4] A diachronic analysis of the corpus identified what we characterized as the "one small problem" caveat running throughout the grant solicitations, whereby affirmation of progress in military technologies is qualified by acknowledgement of outstanding problems, which becomes the justification for additional investments in research. A closer analysis of a subset of Defense Advanced Research Projects Agency (DARPA) calls for the use of AI in battlefield applications further demonstrated the two-way traffic between research communities (both academic and commercial) and the DoD. These showed more specifically how the DoD strategically mobilizes the category of "basic research" to reflect the interests and directions of the research community, which in turn inform the problem formulations and imagined solutions of DoD battlefield "applications." Taken together, we argued that grant solicitations work as a vehicle for the mutual enlistment of DoD funding agencies and the academic AI research community in setting research directions. Bringing together the military trope of enlistment with the concept of enrollment as it has been developed within the field of STS (Callon, 1986), we see mutual enlistment as one among a set of processes through which problems are defined, solutions put forward, and actors recruited into positions and relations that work to sustain a particular definition of their collective situation (in this case, the necessity of AI for military supremacy, and of military supremacy as a basis for U.S. national security).

## 2. Agonistic engagements

Institutionally sanctioned spaces for interdisciplinarity are conventionally opened in response to a problem deemed to require multiple forms of expertise. Those spaces are also, and arguably more importantly, "a means of generating questions around which new forms of thought and experimental practice can coalesce" (Barry and Born, 2013: 10). In this respect, there is a difference that matters between normative research enlisted in the service of agendas where the problem framing is not itself open to question, and research that questions the frame.

The initial ground for our collaboration was a shared desire to trouble the unquestioned assumptions and the political and economic arrangements that frame the purposes and practices of the field of AI. Lucy's engagement with AI began in the mid-1980s (a time of ascending investments in symbolic AI and knowledge representation) as a PhD student in anthropology who found a "field site" for her dissertation at Xerox's Palo Alto Research Center (PARC). Engagement with her colleagues required the acquisition not of comparable technical competency (clearly beyond her reach), but rather the development of a sufficient level of conceptual understanding to be able to identify issues at the intersection of her interests and those of her colleagues and to make her reading of those issues intelligible to them. The first of her collaborations was a critical engagement with the development of an "expert system" based on the then-dominant planning model of AI. Differences between her

[4]Agre (1997a) describes how early DoD funding programs for academic research in AI set out research areas like problem-solving, learning, vision and planning. Anticipating our analysis, Agre observes that within these broad rubrics, DoD funding agencies could articulate the relevance of the research to their strategic objectives while also providing faculty and students with a sense of research autonomy. In his own case, Agre attributes the positive reception of his work with David Chapman on models of situated activity (Agre and Chapman, 1990) to a shift in military strategy toward autonomous battlefield robots under the Strategic Computing Initiative. For a critique of that initiative written at that time, see (Ornstein et al., 1984).

https://doi.org/10.1017/cfc.2025.10004 Published online by Cambridge University Press

own interactionist analysis of sociality (drawn from developments at the time in anthropology, ethnomethodology, and Conversation Analysis) and that of her colleagues provided fertile grounds for a critical reframing of the problem of human-machine communication (Suchman 1987, 2007).

Beginning his career at IBM Watson in 2015, Sireesh imagined his role as a developer of NLP systems in industry to be about making useful things.[5] What he encountered was a gap between the project of advancing NLP research and development on one hand, and the actual requirements of practical use on the other. He found himself building systems that were necessarily ad hoc and aimed at the automation of white-collar office work, an aim that aligned only loosely with the affordances of the NLP technology of the moment. Moving in 2018 to a self-identified AI startup, Sireesh was hired to develop NLP to support financial analysis. He soon discovered that the company was positioned more to be acquired than to produce working technologies: the purpose of the systems that he developed was to create impressive prototypes in which NLP was more of a promise than a fully implemented working system. At the end of his time there, he was team lead for machine learning operations, charged with engineering ongoing assessments of how models were performing, facilitating new data collection, and retraining models when required.

The divergence of new data from the data a model was trained on, however – both at IBM and at the startup – required labour-intensive re-annotation by people with relevant expertise. Overall, the mandate to deliver even partial workplace automation through AI leads, in Sireesh's observation, to progressive immiseration in the workplace, as the work becomes annotating data for the machine rather than doing the job one is hired for. Billed as transformative, contemporary AI reinforces the promise to management of a comprehensive representation of business processes, independent of the actual labor involved. Traditional computer programming models process in code, depending on the user or operator to enact the parts of the process that cannot be encoded. Contemporary data-driven AI encodes process in weights, which expand the range of (especially imprecise and underspecified behavior) that can be modeled, while at the same time making understanding and editing of the modeled behavior much more difficult. In the end, the job modelled is a further devalued rendering of the job as it existed before, where the view of the work process that practitioners used to have has been encapsulated in a model that reflects none of the nuance with which they understand their practice, and whose logics are even more opaque. What is left to workers, then, is to correct the deficient model in the hopes of recovering the job that they performed, through the language that the model understands: annotated data.

While at the startup, Sireesh proposed an audit for gender bias in a speech-to-text algorithm trained on a severely skewed dataset of earnings call transcripts. These calls consist largely of audio from C-suite executives and analysts discussing public companies' quarterly performance; as of 2023, an S&P Global report estimates that only 11.8% of C-suite positions in publicly traded U.S. companies are held by women (Chiang et al., 2024). Sireesh's manager made it clear that even if the audit were to uncover issues, they would not be acted upon unless they directly contributed to improving the model's metrics on the same skewed dataset. Pursuing his concern with bias in NLP work, Sireesh decided to return to academia, still imagining his project as a technical one but with a labour-focused view aimed at giving agency back to users. Approaching the field of AI skeptically, with interests arising from his experience of failed projects in industry rather than from historically given problems in computer science, left Sireesh with the question: Where is the space to build software in thoughtful ways within the existing economic system? His move into the academy was motivated by the possibility that he might pursue this question as one integral to his profession, rather than being professionally complicit while engaging in critique on the side. In the academy again, however, he found that the concerns shaping NLP research and development are subject to industry pressures. The field is also

[5]This was at the outset of a push to commercialize NLP technologies following the widely celebrated release of the Watson Jeopardy system (see https://www.ibm.com/history/watson-jeopardy). As Sireesh observed, the connection between Watson and what developers ended up building for commercial clients was a loose one, as the full Watson system was custom-engineered for Jeopardy in a way that somewhat precluded its use in other arenas.

https://doi.org/10.1017/cfc.2025.10004 Published online by Cambridge University Press

narrowly circumscribed by technical mandates that marginalize the orientation to usefulness, while initiatives such as "NLP for good" and audits for bias operate as a separate and peripheral thread of NLP research.

Coming from a liberal arts background but with a lifelong love of computers, David undertook a PhD in software engineering (a field considered interdisciplinary within the rubric of academic computer science). During his studentship, he was hired at Microsoft Research (MSR) to conduct interviews examining why engineers were unable to use the responsible AI guidance provided to them in the form of a booklet of procedures. Expected to return results along the lines of "The font size is too small here," David instead asked further questions more akin to, "Do you feel supported by your manager in acting on this guidance?" These questions were judged to introduce methodologically unacceptable "bias" beyond the set scope of the project, and he was asked to train further in how to conduct unbiased and ethical interviews. Rather than disciplining his qualitative research methods, however, David's further reading led him to shift from viewing interviewing as a method to a deeper theoretical analysis of the interview as an encounter, and to a wider understanding of the differing epistemological stances that comprise qualitative research (Guba and Lincoln, 1994). Like Sireesh, he experienced how, entrained in the prevailing instrumental view of interdisciplinarity in the service of the AI project, initiatives that question the fundamental premises or goals of the project are systematically sidelined. This problem is clear in the operationalization of ethics in AI in the form of a checklist or guideline that serves as a substitute for thought and critical engagement.

In contrast to the premise that there are given problems that require interdisciplinarity for their solution, in sum, our shared experience of interdisciplinarity is as a pathway to questioning the frame that has been handed to us in order to reformulate the problem. This is, necessarily, what Barry and Born (2013: 12) identify as an "agonistic encounter," in these cases disclosing the many ways in which computer science has encapsulated itself, only expanding its boundaries in instrumental ways. As a consequence, the aspects of the discipline that engage with the world outside the closed theoretical/technical domain of academic computer science (e.g., interface design, software engineering) are marginalized. While these subfields are necessarily more central to practices of industrial research and development, the latter reinforce the commitment to closure over problematization in the interest of the delivery of products. In the current moment of rapid expansion and commercialization of AI-enabled technologies, moreover, opening the field of AI requires radical modes of disciplinary reconfiguration that question not only received assumptions and methodological commitments, but also the political economies and professional subjectivities through which technical practitioners are disciplined.

## 3. Critical technical practice

In "Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI," Philip Agre (1997a) points to broader frictions between computer science and what that field considers to be its cognate disciplines. "In its representational aspirations," he writes, "computing has been constituted as a kind of imperialism; it aims to reinvent virtually every other site of practice in its own image" (Agre, 1997a: 131). Agre contextualizes the field of AI within the broader cybernetic turn and the intensified entanglements of computer science with the military-industrial complex emerging from World War II and the Cold War. He identifies two related precepts that underwrite the AI project: that cognition comprises mental operations translatable as computation, and that those translations can be achieved through techniques of formalization. These precepts are sustained, he argues, rhetorically: "As a practical matter," Agre observes, "the purpose of AI is to build computer systems whose operation can be narrated using intentional vocabulary" (1997a: 139). The dominant narrative for that project in the 1980s was the so-called planning model, in which plans – understood as mental representations – provide the executable instructions for action.

https://doi.org/10.1017/cfc.2025.10004 Published online by Cambridge University Press

As PhD students during the 1980s at Xerox PARC and MIT respectively, Lucy and Phil Agre met in the context of their shared critical engagement with the planning model. Conceptual and methodological resources from anthropology and ethnomethodology, combined with a technically informed understanding of the prospects and limits of AI planning, provided fertile ground for a critique of the representation of plans as mental constructs determining action. We argued instead for an articulation of plans as discursive and material artifacts, generated within and mobilized as resources for situated activity. Along with prior scholarship in the humanities and social sciences (e.g. Dreyfus, 1992; Forsythe, 1993) our critiques drew from a small but growing body of work within computer science concerned with AI's foundations and ramifications – not only as a matter of effects outside the discipline, but also for its constitutive technical practices (e.g. McDermott, 1976; Weizenbaum, 1976).[6]

In *Computation and Human Experience* (1997b), Agre formulates the concept of a critical technical practice as a stance open to its own reconfiguration. This requires, he argues, a rejection of the premise of disciplinary foundations as an underlying structure, in favour of attention to a discipline's workings as a historically specific practice. Integral to generative disciplinary practice, on this premise, is continuing critical inquiry into that history. Rather than seeing the associated risk to disciplinary concepts and methods as a threat, he proposes that such risk comprises "the promise of better ways of doing things" (1997b: 23). Embracing this stance, Agre maintained his conviction that AI as a field could contribute to a deeper understanding of human experience, but only insofar as it developed a reflexive critique of the genealogy and consequences of its technical language. His own voracious and extensive reading beyond the discipline was, for him, a way "to extricate myself from the whole tacit system of intellectual procedures in which I had become enmeshed during my years as a student of computer science" (1997b: 148). Central to that process for Agre was a resistance to what he identified as the false precision of technical formalization (1997a: 12).

Developments in the field of AI since Agre's interventions are mixed. Consistent with his call is a growing body of critical analysis coming from the various disciplines that contribute to the current AI project, including computational linguistics/natural language processing and machine learning (e.g. Bender et al., 2021; Gururaja et al., 2023), along with scholarship in the interdisciplinary fields of information studies and STS (e.g Burrell, 2016; Ribes et al., 2019; Seaver, 2017). These are joined by activist collectives and progressive research institutes within commercial and academic computing.[7] These critiques question the validity of narrow benchmarks as measures of the quality of AI systems, insisting that adequate assessment of systems can only be done in the context of their applications and consequential effects. At the same time, the massive growth of (largely speculative) investment in AI has exploited the evocative vagueness of the language of the field to expand beyond academic confines, shape public opinion, and promote uncritical narratives of technological solutionism and inevitability. As the rhetorical and financial reach of those with vested interests expands, the work of critical resistance requires wider coalitions both within and beyond the academy.

## 4. Auspicious interdisciplinarity

For critical technical practitioners, moving outside the disciplinary boundaries of computer science affords intellectual resources for understanding how those boundaries are constituted and with what effects, and how they might be redrawn. While arguably interdisciplinary from its beginnings, at least in terms of those disciplines that make up the field of Cognitive Science (principally cognitive

[6]For a recent reflection on the experience of "awakening" to critical technical practice, see (Malik and Malik, 2021).

[7]For example, the Tech Workers Coalition (https://techworkerscoalition.org/), No Tech for Apartheid (https://www.notechforapartheid.com), Athena (https://athenaforall.org/), as well as the AINow Institute (https://ainowinstitute.org/) and the Distributed AI Research Institute (https://www.dair-institute.org/).

https://doi.org/10.1017/cfc.2025.10004 Published online by Cambridge University Press

psychology, computer science and philosophy), the AI project has been implemented as and through the techniques and technologies of computer science and software engineering, and those technical fields retain their dominance. Recently, however, the boundary-crossing traffic has thickened as critical practitioners are moved to expand and deepen their understanding of the starting premises, limits, and consequences of their technical work, while those with primary affiliations to the humanities and social sciences expand and deepen their technical literacy. "Auspicious interdisciplinarity," Barry and Born propose, "is associated not only with the constitution of new objects, but with the cultivation of interdisciplinary subjectivities and skills" (2013: 39).

Early critiques of the disciplinary limits of AI, then a somewhat esoteric project, focused on the question of whether and how computational models of mind, action, and interaction could do justice to the lived, organically embodied, culturally embedded, and historically evaluated relations that comprise intelligence. Now, more than ever, those disciplinary critiques require a radically extended frame—one that includes the political economies of what Hao (2025) has identified as the "Empire of AI" (see also Crawford, 2021). That reframing insists on attention to the conditions of possibility for the scaling of AI, from resource extraction to massive new investments in energy production, and from exploited labour to concentrations of wealth and associated power. In this moment of frantic investment and expanding colonization by governments and interested corporations, more and more of us are enrolled as (often unwitting) contributors to the AI project through the capture of our data traces, as outsourced labour, or as ingrained consumers. Adequate critique in the present moment calls for a critical technical practice committed to asking fundamental questions, resisting the posited inevitability of AI, and generating more radical alternatives that reject technosolutionism in favor of collective movements to build just and sustainable futures. Whether and how algorithmically based technologies could enable those futures can only be determined through continually unfolding experiments in auspicious interdisciplinarity.

**Funding statement.** None.

**Competing interests.** The authors declare none.

## References

**Adam, A.** (1998) *Artificial Knowing: Gender and the Thinking Machine*. New York: Routledge.

**Agre, P.** (1997a) Toward a critical technical practice: Lessons learned in trying to reform AI. In G. Bowker, L. Gasser, S.L. Star and B. Turner (Eds.), *Social Science, Technical Systems, and Cooperative Work: Beyond the Great Divide* (pp. 131–57). Erlbaum.

**Agre, P.** (1997b) *Computation and Human Experience*. New York: Cambridge University Press.

**Agre, P. and D. Chapman** (1990) What are plans for? *Robotics and Autonomous Systems*, *6*(1–2), 17–34. https://doi.org/10.1016/S0921-8890(05)80026-0

**Barry, A., and G. Born**, (Eds.) (2013) Introduction. In *Interdisciplinarity: Reconfigurations of the Social and Natural Sciences* (pp. 1–56). London: Routledge.

**Bender, E., T. Gebru, A. McMillan-Major and S. Shmitchell** (2021) On the dangers of stochastic parrots: Can language models be too big? *FAccT 21*. https://doi.org/10.1145/3442188.3445922

**Burrell, J.** (2016) How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data and Society*, *3*(1), 1–12. https://doi.org/10.1177/2053951715622512

**Callon, M.** (1986) Some elements of a sociology of translation: Domestication of the scallops and the fishermen of St. Brieuc Bay. In J. Law (Eds.), *Power, Action and Belief - A New Sociology of Knowledge?* (pp. 196–233). London: Routledge.

**Chiang, H., S. Kaulapure and D. Sandberg** (2024) *Elusive Parity: Key Gender Parity Metric Falls for First Time in 2 Decades*. S&P Global Market Intelligence Quantamental Report, March. https://doi.org/10.2139/ssrn.4877730

**Crawford, K.** (2021) *Atlas of AI*. New Haven: Yale University Press.

**Dreyfus, H.** (1992) *What Computers Still Can't Do*. Cambridge, MA: MIT Press.

**Forsythe, D.** (1993) Engineering Knowledge: The Construction of Knowledge in Artificial Intelligence. *Social Studies of Science*, *23*, 445–77. https://doi.org/10.1177/0306312793023003002

**Guba, E.G. and Y.S. Lincoln**. (1994) Competing paradigms in qualitative research. *Handbook of Qualitative Research*, *2*(163–194), 105.

https://doi.org/10.1017/cfc.2025.10004 Published online by Cambridge University Press

**Gururaja, S., A. Bertsch, C. Na, D.G. Widder and E. Strubell** (2023) To build our future, we must know our past: Contextualizing paradigm shifts in natural language processing. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 13310–25, Singapore. Association for Computational Linguistics.

**Hao, K.** (2025) *Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI*. Penquin.

**Jones, M.** (2016) *Reckoning with Matter: Calculating Machines, Innovation, and Thinking about Thinking from Pascal to Babbage*. Chicago University of Chicago Press.

**Mackenzie, A.** (2017) *Machine Learners: Archaeology of a Data Practice*. Cambridge, MA: Cambridge University Press.

**Malik, M. and M. M. Malik** (2021) Critical technical awakenings, in. *Journal of Social Computing*, *2*(4), 365–84. December. 2021. https://doi.org/10.23919/JSC.2021.0035

**McDermott, D.** (1976) Artificial Intelligence Meets Natural Stupidity. *SIGART Newsletter*, (57), 4–9. https://doi.org/10.1145/1045339.1045340

**Ornstein, S., B. C. Smith and L. Suchman** (1984) Strategic Computing. *Bulletin of the Atomic Scientists*, *40*(10), 11–15. https://doi.org/10.1080/00963402.1984.11459292

**Pasquinelli, M.** (2023) *The Eye of the Master: A Social History of Artificial Intelligence*. London and New York: Verso

**Ribes, D., A. S. Hoffman, S. Slota and G. Bowker** (2019) The logic of domains. *Social Studies of Science*, *49*(3), 281–309. https://doi.org/10.1177/0306312719849709

**Riskin, J.** (Eds.) (2007) *Genesis Redux: Essays in the History and Philosophy of Artificial Life*. Chicago: University of Chicago

**Seaver, N.** (2017) Algorithms as culture: Some tactics for the ethnography of algorithmic systems. *Big Data and Society*, *4*(2), 1–12. https://doi.org/10.1177/2053951717738104

**Suchman, L.** (1987) *Plans and Situated Actions: The problem of human-machine communication*. Cambridge: Cambridge University Press.

**Suchman, L.** (2007) *Human-Machine Reconfigurations*. Cambridge: Cambridge University Press.

**Weizenbaum, J.** (1976) *Computer Power and Human Reason*. New York: W.H. Freeman.

**Widder, D.G., Gururaja, S., and Suchman, L.** (2024) Basic Research, Lethal Effects: Military AI Research Funding as Enlistment. arXiv:2411.17840

**Widder, D. G., and Nafus, D.** (2023) *Dislocated accountabilities in the "AI supply chain": Modularity and developers' notions of responsibility*. Big Data & Society, 1–12.

**Widder, D. G., West, S. M., and Whittaker, M.** (2024) Why 'open' AI systems are actually closed, and why this matters. *Nature*, 635, 827–833.

**Lucy Suchman** is Professor Emerita of the Anthropology of Science and Technology at Lancaster University in the UK.

**Sireesh Gururaja** is a PhD student at Carnegie Mellon University's Language Technologies Institute.

**David Gray Widder** is a Postdoctoral Fellow at the Digital Life Initiative at Cornell Tech, and earned his PhD from the School of Computer Science at Carnegie Mellon University.

**Cite this article:** Suchman, L., Gururaja, S. and Widder, D.G. 2025. AI interdisciplinarity as critical technical practice. *Cambridge Forum on AI: Culture and Society*, **1**, e2, 1–8. https://doi.org/10.1017/cfc.2025.10004

https://doi.org/10.1017/cfc.2025.10004 Published online by Cambridge University Press