

RESEARCH ARTICLE

Animo nullius: on AI's origin story and a data colonial doctrine of discovery

Jonnie Penn

Department of History and Philosophy of Science, University of Cambridge, UK
Email: jnp28@cam.ac.uk

Abstract

This paper traces elements of the theoretical origins of artificial intelligence to capitalism, not neurophysiology. It considers efforts in the twentieth and twenty-first centuries to formalize a science of mental behaviour using the dynamics of social rather than neural phenomena. I first revisit early American theorists' controversial ambivalence toward neurophysiology, showing how this group benefited from post-war corporate and military investments in commercial and imperial expansion, which sustained and expanded their influence over the emerging field. I then trace the lasting effect of the founders' early rhetoric through AI's institutionalization after 1960, arguing that from the 2010s technology corporations set out to veil their enclosure of the data commons via appeal to a curious precedent: the scientific pedigree of AI. By relating the field to the history of capitalism, and specifically the rise of assetization in modern technoscience, I invite reflection on AI's origin story and on broader parallels between historical colonialism and data colonialism. I offer a heuristic – *animo nullius*, for 'no persons' mind' – as an attempt to name rhetorical manoeuvres that leverage the authority of mind-as-computer metaphors in order to naturalize acts of seizure.

You cannot discover an inhabited land. Otherwise I could cross the Atlantic and 'discover' England.

Dehatkadons, a traditional chief of the Onondaga Iroquois

In his 1986 book *The Society of Mind*, Marvin Minsky outlined a theory of intelligence informed by three decades of research on artificial intelligence, a field he helped found. Minsky stated his aspiration to emulate the foundational insights of Galileo and Newton by reducing the dynamics of the mind down to the simplest possible terms. He theorized that intelligence was the product of many simple mental processes, or modular 'agents', joined by cross-connections that could be modelled mathematically via the logic of, in his words, a 'society'.¹ Minsky did not specify which type of society, nor from where within that social structure its all-important dynamics were to be judged. For him, a 'society' was a philosophical abstraction, an idea that bore universally as if its logics were self-evident, meaningful and constant.

1 Marvin Minsky, *The Society of Mind*, New York: Simon & Schuster, 1988, p. 1.

Minsky was neither the first nor the last in his field to formalize the science of mental behaviour using the vocabulary and dynamics of social phenomena.² Other prominent researchers turned to social logics to render neural complexities in ways they deemed susceptible to mathematical treatment. Herbert Simon, a driver of the symbolic paradigm, borrowed from the study of administration on the assumption that decision making in large human organizations had gone unchanged for ‘at least four thousand years’.³ Frank Rosenblatt, a driver of the connectionist paradigm, built his perceptron models on theory borrowed from Hayekian economics. What this and other relevant metaphors had in common was the motivating presumption that the mind was an orderly thing; that it lived inside an individual’s brain; and that it followed an implicit, reliable ‘logic’ that could be convincingly modelled with modes of computation derived from the observation of social events.

These instances from the earliest days of AI point to a broad category of non-neural metaphors that go largely unacknowledged in existing genealogies of the overlap between digital computer programming and neural activity. This omission is surprising given that contemporary machine-learning researchers still turn to subjective social logics to model neural phenomena, including by benchmarking progress on machine intelligence against abstract strategy games (such as chess, go, StarCraft) and by modelling software ‘agents’ into social formations premised on theories of adversarial human relationships (such as game theory, general adversarial networks).⁴

At first glance, the simulation of subjective social logics and their subsequent deployment into knowledge management infrastructures reads as baldly political. This is especially true when such applications are dressed, via selective metaphors, as neural, or framed, via selective histories, as apolitical and apparently self-evident and universal. That the intentions of state and industrial actors who fund(ed) such research often remain unacknowledged adds another level of contingency for historians to unpack. Sensitive to the overlap between social and computational logics in the prehistory of digital computing, Agar advocates reading the history of information technologies and modern state formation side by side. ‘Several of the most important moments in the history of information technology revolve, rather curiously, around attempts to capture, reform, or redirect governmental action’, he argues.⁵

With Agar’s insight in mind, this paper trials a heuristic – *animo nullius*, for ‘no persons’ mind’ – to suggest that uncritical histories of computer-simulated ‘intelligence’ and ‘learning’ advance acts of seizure using metaphor. In choosing this framing, I invite reflection on parallels between historical colonialism and data colonialism, specifically parallels with *terra nullius*, the Roman legal expression for ‘land that belongs to no one’, which remains the spurious philosophical basis for a swathe of contemporary international law. *Terra nullius* – the concept, if not the name – was repurposed opportunistically after the fifteenth century as a rescripting ‘principle’, ‘doctrine’, ‘policy’ and legal ‘fiction’ by other European colonial empires to dispossess Indigenous peoples of their land.⁶ Herein, I question how the authority of foundational language about machine ‘intelligence’ developed, often ambivalently, in the mid-twentieth century is now being

2 Jonnie Penn, ‘Inventing intelligence: on the history of complex information processing and artificial intelligence in the United States in the mid-twentieth century’, PhD dissertation, University of Cambridge, 2020.

3 Herbert A. Simon, *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organizations*, 4th edn, New York: Free Press, 1997, pp. 1–2.

4 For a recent exception see Michael Castelle, ‘The social lives of generative adversarial networks’, *Conference on Fairness, Accountability, and Transparency*, Barcelona and New York: ACM, 2020.

5 Jon Agar, *The Government Machine: A Revolutionary History of the Computer*, Cambridge, MA: MIT Press, 2003, p. 7.

6 Yogi Hale Hendlin, ‘From terra nullius to terra communis’, *Environmental Philosophy* (2014) 11(2), pp. 141–74.

repurposed to normalize enclosure of the data commons. ‘Just as historical colonialism over the long run provided the essential preconditions for the emergence of industrial capitalism’, write Mejias and Couldry, ‘we can expect that data colonialism will provide the preconditions for a new stage of capitalism that as yet we can barely imagine, but for which the appropriation of human life through data will be central’.⁷

This paper proceeds in two parts. First, I recount how the founders of ‘AI’, broadly construed, proselytized that mathematics, logic, computation and statistical formalisms would render mental behaviour legible and amenable to simulation. Research in the mid- to late 1950s on what later became ‘AI’ positioned the human mind as universal, programmable and knowable, a ‘domain’ to settle through the abstract instrumentation of positivist Western technoscience. Figures in symbolic AI like Marvin Minsky, John McCarthy, Herbert Simon and Allen Newell, and in machine learning like Frank Rosenblatt, invoked the implied authority of neural metaphors, as if the correspondence was literal, even while they oscillated over the degree to which neurophysiology should or could serve as the empirical basis of their work. I explore how ties to US-based state-spirited institutions in this period helped them to advance this research.

The second section traces the lasting effect of these rhetorical moves through AI’s institutionalization in the twentieth and twenty-first centuries. I show how large American technology corporations repurposed this ambiguously ‘neural’ rhetoric in the 2010s and 2020s to build and expand markets for the sale of proprietary software and file-hosting services. In the present day it is global corporations, rather than upstart American researchers, that equivocate over the authority (and, with it, the historical genealogy) that underwrites AI. Hanging in the balance are emerging societal norms about the degree to which ‘AI’ outputs are provably transformative and protectable via copyright or other law. As with *terra nullius*, both those who make and those who contest such claims will look to history for adjudication. I call attention to *animo nullius* to invite historians of technology to pay closer attention to the ways in which AI’s theoretical origins built on the edifice and practices of capitalism more substantially than on neurophysiology, even when neural metaphors were applied.

The origins of AI

In the mid-1950s, researchers in the United States melded formal theories of problem solving and intelligence with another powerful new tool for control: the electronic digital computer. Several branches of Western mathematical science emerged from this nexus, including computer science (1960s–), data science (1990s–) and artificial intelligence. Nascent efforts to formalize the dynamics of human thought elaborated what Roberto Cordeschi has called ‘the culture of the artificial’, in which mental processes were seen as independent of organic structures and could be verified as such.⁸ During the first half of the twentieth century, this culture moved from the fringes of psychology to mainstream research on cognitive science.

In the 1950s, this tradition converged with the advent of digital electronic computing. Between 1952 and 1957, new computational mechanisms like ‘assemblers’ and ‘compilers’ consolidated the prohibitively laborious instruction sets used to ‘program’ a digital computer into ever more accessible and coherent programming ‘languages’. In the US, industrial actors like IBM advanced and leveraged these consolidations to usher new paying

⁷ Nick Couldry and Ulises A. Mejias, ‘Data colonialism: rethinking big data’s relation to the contemporary subject’, *Television & New Media*, 2 September 2018, pp. 336–49, 337.

⁸ Roberto Cordeschi, ‘The discovery of the artificial: some protocybernetic developments 1930–1940’, *AI & Society* (July 1991) 5(3), pp. 218–38, 218.

audiences into standardized modes of computing. In 1958, the *Communications of the Association for Computing Machinery*, a new journal, fielded letters from an emerging class of professionals and theorists over what to name their profession. Suggestions included ‘synnoetics’ (‘science of the mind’ in Greek), ‘computer science’ and the generalizable notion of ‘comptology’ (e.g. nuclear comptologist, logistics comptologist).⁹

Freighted in these terms were deep claims about the explanatory potential of information-processing metaphors to describe neural phenomena, claims that often attracted explicit attention at first, only to be silently absorbed over time.¹⁰ Elsewhere I have shown that despite their differing foundational approaches to AI and machine learning, key researchers blurred the line between metaphorical and literal descriptions of the intersections between computation and cognition.¹¹ Anthropomorphic language was immediately controversial. The 1953 book *Automatic Digital Calculators* apologized for using ‘memory’ in relation to computing and acknowledged that terminology in the field had ‘not yet stabilized’.¹² Prominent American engineers and mathematicians like Claude Shannon and John von Neumann bristled at a new generation’s hubristic use of mind-as-computer metaphors and took pains to articulate their limits. Shannon dismissed John McCarthy’s attempt to include ‘intelligence’ as a relevant descriptor in *Automata Studies*, their co-edited 1956 volume.¹³ At the 1948 Hixon Symposium on Cerebral Mechanisms in Behavior, von Neumann positioned his mathematical work on neural phenomena as that of an ‘outsider’ prohibited from meaningful speculation on anything other than the dynamics of idealized versions of elementary physiological units.¹⁴ In the early 1950s, professional etiquette compelled Marvin Minsky to use scare quotes around his early mathematical notions of machine ‘learning’.¹⁵

By the late 1950s, these quotes had disappeared. Machine ‘learning’ became machine learning. To assert the perceived legitimacy of this transition from metaphorical to vaguely literal descriptions of human capacities in machines, these men permitted clannishness, self-aggrandizement, speculative rhetoric, fluid definitions of key terms and poor citation practices, actions that drew attention to questions of how to accomplish such aims and away from whether they were well founded.¹⁶ Even as they disagreed over the correct measure of fidelity to biological phenomena in their modelling and the appropriate rhetoric for their results, Simon, Rosenblatt, McCarthy and Minsky positioned brain modelling as ripe for a grand theory of cognition, a possibility they all presumed existed. ‘Once one system of epistemology is programmed and works no other will be taken seriously unless it also leads to intelligent programs’, McCarthy wrote in his diaries. ‘The artificial intelligence problem will settle the main problems of epistemology in a

9 Quentin Correll, Otto Matty Khodr and Alexander Vanderburgh, ‘Letters to the editor’, *Communications of the ACM* (1 July 1958) 1(7), pp. 2–3.

10 Matthew Cobb, *The Idea of the Brain: The Past and Future of Neuroscience*, New York: Basic Books, 2020; Penn, op. cit. (2).

11 Penn, op. cit. (2).

12 Kathleen H.V. Booth and Andrew D. Booth, *Automatic Digital Calculators*, London: Butterworth Scientific Publications, 1953, p. v.

13 Ronald Kline, ‘Cybernetics, automata studies, and the Dartmouth conference on artificial intelligence’, *IEEE Annals of the History of Computing* (April 2011) 33(4), pp. 5–16, 8; Kline, *The Cybernetics Moment: Or Why We Call Our Age the Information Age*, Baltimore: Johns Hopkins University Press, 2015, p. 157.

14 John von Neumann and Lloyd A. Jeffress, ‘The general and logical theory of automata’, in Lloyd A. Jeffress (ed.), *Cerebral Mechanisms in Behavior*, London: Chapman & Hall, 1951, pp. 1–41.

15 Marvin Minsky, ‘Theory of neural-analog reinforcement systems and its application to the brain-model problem’, Princeton University, 1954, photocopy, Ann Arbor, Michigan University Microfilms, 1962, pp. 3–2.

16 See Penn, op. cit. (2), pp. 64; James Fleck, ‘Development and establishment in artificial intelligence’, in Norbert Elias, Herminio Martins and Richard Whitley (eds.), *Scientific Establishments and Hierarchies*, Dordrecht: Springer Netherlands, 1982, pp. 169–217.

scientific way.¹⁷ Even Rosenblatt, who called AI researchers his ‘loyal opposition’ because their deterministic ‘paper exercises’ contravened biological evidence, neglected to state the limits of his own instrumentalist brand of statistical modelling. The difference between his commitments to biological fidelity and those of his ‘opposition’ was one of degree, not of kind.

Internalist accounts of early brain model theory have tended to forgive these common commitments as customary. Recurring emphasis on narratives of discontinuity, such as on funding ‘winters’ or the fabled 1956 Dartmouth Summer Research Project, where AI gained its name, have obscured broader methodological and structural continuities.¹⁸ One recurring, yet distracting, tendency in existing histories is a focus on the swing of fashions between two apparent poles: symbolic reasoning (such as the simulation of problem solving) and neural-network research (such as the simulation of learning). I eschew that dichotomy here to revisit what these pursuits ultimately had in common, such as the foundational beliefs that complex mental processes were susceptible to mathematical formalization, that digital electronic computers were the appropriate tool for that job, and that the product of this enterprise would be an abstract ‘language’.

Historians have questioned another commonality: the role that generous institutional patronage played in initial efforts to formalize the mathematics of thought. Kline illustrates that Dartmouth was not the launch pad it has been described as: funding was halved, participants dropped out or left after a day, and the group’s diverse research foci sprawled.¹⁹ That the workshop did not establish AI draws our attention to other funding and organizational dynamics that helped to bolster the field thereafter. Jean Pierre Dupuy considers the fates of two adjacent research areas. He describes cybernetics as a ‘failure’ in comparison to information theory, which found a professional home in the Institute of Electrical and Electronics Engineers.²⁰ Andrew Pickering argues that cybernetics failed because it lacked similar backing. The Macy Conferences and Ratio Club lacked the capacity to formally train students, set a stable research agenda or grant degrees, hampering the ability of cybernetics to propagate, at least initially. He writes,

When we think of interdisciplinarity we usually think of collaborations across departments in the university ... The centre of gravity of cybernetics was not the university at all. Where was it then? The simplest answer is: nowhere ... Cybernetics flourished in the interstices of a hegemonic modernity, largely lacking access to the means of reproduction: the educational system.²¹

This was not the case for AI. Less than a decade after the Dartmouth meeting, the AI group at MIT (and then at Stanford and Carnegie Mellon) operated with an annual budget purported to be in the millions – a remarkably brief transformation into institutionalization

17 John McCarthy, ‘Methodology of work on the artificial intelligence problem’, Stanford, CA, n.d., p. 81, John McCarthy Papers (SC0524); Mathematics, 1955–?, Department of Special Collections and University Archives, Stanford University Libraries.

18 In the most popular textbook on AI, Dartmouth and 1956 are described as the time and place of ‘the birth of artificial intelligence’. See Stuart J. Russell, Peter Norvig and Ernest Davis, *Artificial Intelligence: A Modern Approach*, 3rd edn, Upper Saddle River: Prentice Hall, 2010, p. 17.

19 Kline, *The Cybernetics Moment*, op. cit. (13), p. 163.

20 Jean-Pierre Dupuy, *The Mechanization of the Mind: On the Origins of Cognitive Science*, Princeton, NJ: Princeton University Press, 2000, p. 15.

21 Andrew Pickering, *The Cybernetic Brain: Sketches of Another Future*, Chicago: The University of Chicago Press, 2011, pp. 215–16.

that historians have yet to properly interrogate.²² What I seek to argue here is that alignments between ‘AI’ and corporate prerogatives eased its access to the means of reproduction. My focus in this gestational period of the mid- to late 1950s is on support provisioned through IBM, the Rockefeller Foundation and (albeit to a lesser extent) Bell Laboratories, whose reasons for involvement encompassed more than a search for a grand theory of cognition. Their incentives included the potential for commercial expansion, the standardization of manufacturing techniques and support for American dominance abroad.

The institutionalization of AI

The Rockefeller Foundation, which funded the Dartmouth meeting credited with giving ‘AI’ its name, was in the mid-1950s a ‘parastate’ or “state-spirited” organization sitting at the heart of the emerging East Coast U.S. foreign-policy establishment.²³ Along with other East Coast philanthropic foundations born from the exploits of nineteenth-century corporate titans (such as Ford, Carnegie), the foundation was keen to enlist ‘ivory tower’ academics to secure global peace by promoting American hegemony against the USSR. This overarching strategy goes part of the way to explaining why the foundation elected to fund contentious research. In 1955 Warren Weaver, then director of its Division of Natural Sciences, expressed misgivings when McCarthy approached him to fund the Dartmouth meeting. He directed McCarthy to Robert Morison, head of Rockefeller’s Biological and Medical Research Division, but told Morison privately, ‘I am very doubtful that the RF ought to do anything about this.’²⁴ Morison, too, expressed qualms. ‘This new field of mathematical models for thought ... is still difficult to grasp very clearly.’²⁵

Morison approved a \$7,500 grant, half the requested \$13,500. He questioned whether a seminar format was the right mechanism for discovery since mathematical theories of brain functions remained ‘pre-Newtonian’.²⁶ His hesitation was not purely subject-based: Kline notes that Morison and Weaver had by then co-funded research on mathematical biology undertaken by Wiener and Rosenblueth, which led to the book *Cybernetics*.²⁷ Central to his doubt was the group’s neglect of existing experts and established knowledge practices. Morison urged the organizers to include psychologists like Hans-Lukas Teuber and Karl H. Pribram, ‘if only for the purpose of keeping the group from speculating too wildly on how the brain might work’.²⁸ While psychologists did not ultimately participate, Morison nonetheless agreed to fund the meeting as a ‘modest gamble for exploring a new approach’.²⁹

IBM’s initial participation in ‘AI’ research, broadly construed, was not wholly dissimilar to the Rockefeller Foundation’s. In the decade after the Second World War and in keeping

22 Daniel Crevier, *AI: The Tumultuous History of the Search for Artificial Intelligence*, New York: Basic Books, 1993, p. 65.

23 Inderjeet Parmar, ‘American hegemony, the Rockefeller Foundation, and the rise of academic international relations in the US’, in Nicolas Guilhot (ed.), *The Invention of International Relations Theory*, New York: Columbia University Press, 2011, pp. 182–219, 186.

24 Warren Weaver, ‘Inter-office correspondence, WW to RSM’ (14 June 1955), Record Group 1.0002, Series 200, Box 26, Folder 219, Rockefeller Archive Center, Sleepy Hollow, NY; Kline, *The Cybernetics Moment*, op. cit. (13), p. 80.

25 As cited in Kline, *The Cybernetics Moment*, op. cit. (13), p. 10.

26 Robert Morison, ‘New York’, 17 June 1955, Record Group 1.0002, Series 200, Box 26, Folder 219, Rockefeller Archive Center, Sleepy Hollow, NY.

27 Even with this support, Kline, *The Cybernetics Moment*, op. cit. (13), pp. 7, 161, like Pickering, positions cybernetics as having succumbed to ‘disunity’.

28 Morison, op. cit. (26).

29 Robert Morison, ‘Letter from Robert Morison to John McCarthy’, 30 November 1955, Record Group 1.0002, Series 200, Box 26, Folder 219, Rockefeller Archive Center, Sleepy Hollow, NY.

with its name, IBM exercised ambitions to broaden its reach internationally. From 1914 until his death in 1956, the company's leader, Thomas Watson Sr, a uniquely talented salesman, pursued growth mercilessly. Echoing Adam Smith's view that international trade was an instrument by which to spread civility and shared prosperity abroad, Watson emblazoned the slogan 'World Peace through World Trade' on IBM's New York headquarters. Paju and Haigh describe the company in the post-war reconstruction period as having operated in a distinctly imperial fashion, planting flags on prominent buildings in major cities, erecting and retaking control over national subsidiaries (ensuring that none were so self-sufficient as to defect or be annexed), and pursuing trade-heavy 'inter-change' manufacturing across Europe.³⁰

The commercial case for automatic coding – and, with it, proto-AI – had not been obvious to IBM leadership. Nathaniel Rochester, an employee and co-convenor of the Dartmouth workshop, claimed in personal records that he and his team had 'sold the idea' of automatic coding to management as a patriotic act after IBM reactivated its military product unit in response to the Korean War.³¹ Having been hired by the navy in 1947 to construct the arithmetic unit of Whirlwind I, amongst the first digital computers in the US, Rochester was particularly well positioned to recognize the prospects of the field. He spent the rest of his career at IBM and recalled his initial surprise to learn that the company had 'no intention of venturing into stored-program machines'.³² Once approved, responsibility fell to him to design and construct their first large-scale scientific computer, the Defense Calculator, released in 1952 as the IBM 701. By 1954, he oversaw a staff of 450 people.³³ After 1956, when Thomas Watson Jr became CEO, IBM's involvement in digital electronic computing deepened significantly as adoption of the technology expanded.

Support from IBM in the mid- to late 1950s became important to the broader AI project. In February 1955, McCarthy left a brief teaching position at Stanford to become assistant professor of mathematics at Dartmouth College. Despite a long-standing interest in computing, he had never actually tried to program a computer until that year.³⁴ This was partly beyond his control; neither Stanford, where he had been based, nor Dartmouth, at which he had just arrived, owned a stored-program computer for him to experiment with. To his good fortune, Philip M. Morse, the 'founding father of operations research', had just convinced MIT to open a new three-storey IBM-funded Computation Center to twenty-five sister colleges and universities across New England.³⁵

As the chosen representative of Dartmouth College, McCarthy was introduced to Rochester, who invited him to spend the summer of 1955 at IBM. A press release from December 1956 described the centre as the 'largest and most versatile data processing facility yet to be made available primarily for education and basic research'.³⁶ As

30 Petri Paju and Thomas Haigh, 'IBM rebuilds Europe: the curious case of the transnational typewriter', *Enterprise & Society* (2016) 17(2), pp. 265–300, 266.

31 Andrew Goldstein, 'Oral history: Nathaniel Rochester, 6 June 1991, interview # 080 for the IEEE History Center', the Institute of Electrical and Electronics Engineers, Inc.; Gerald W. Brock, *The Second Information Revolution*, Cambridge, MA: Harvard University Press, 2003, pp. 98–9; On IBM's decision see Paul N. Edwards, *The Closed World: Computers and the Politics of Discourse in Cold War America*, Cambridge, MA and London: MIT Press, 1997, p. 381.

32 F.J. Gruenberger, 'The history of the JOHNNIAC', *Annals of the History of Computing* (July 1979) 1(1), pp. 49–64, 50.

33 Nathaniel Rochester, 'Biographical data, Nov 3, 1957', Nathaniel Rochester Papers, Box 1, Folder 6, Biographical File 1956–60, Library of Congress, Washington, DC (accessed 1 March 2017), p. 3.

34 Nils J. Nilsson, 'John McCarthy, 1927–2011', National Academy of Sciences, 2012, p. 5, at www.nasonline.org/publications/biographical-memoirs/memoir-pdfs/mccarthy-john.pdf (accessed 28 October 2023).

35 John D.C. Little, 'Philip M. Morse and the beginnings', *Operations Research* (February 2002) 50(1), pp. 146–8.

36 Atsushi Akera, *Calculating a Natural World: Scientists, Engineers, and Computers during the Rise of U.S. Cold War Research*, Cambridge, MA: MIT Press, 2007, p. 287.

Dartmouth's appointee, McCarthy was invited to make use of the 18,000-square-foot facility and twenty-five-member staff.³⁷ As part of this arrangement, McCarthy was asked to 'sell the idea' of computing to his peers in academia. 'I think that sooner or later programming will become as basic a part of a scientific education as calculus, and strongly advise learning it even if you don't see immediate application to your problems', he wrote.³⁸ New England academics were receptive; 40 per cent of his students under the IBM flagship were professors. Access to the centre helped McCarthy deepen his collaborations with Minsky. In September 1957, he became an MIT Sloan fellow in the physical sciences, a position that allowed him to relocate to the centre full-time. In September 1958, he and Minsky co-founded the MIT Artificial Intelligence Project with two programmers, a secretary, a typewriting machine and six graduate students.

The search for a theory of 'intelligence' – and a computational language to express it – was expensive, but alignment with state and corporate institutions, and their motivating logics, helped settle the bill and advance the project into a stand-alone discipline. The two cases just outlined speak to a general alignment between mind-as-computer research and state-industrial aims. Of the twenty or so core participants at the Dartmouth workshop, seven were affiliated with IBM, three with Bell Laboratories, two with the RAND Corporation and twelve with MIT, one of America's largest non-industrial defence contractors at the time. IBM, Bell Laboratories and other home organizations paid staff salaries to attend. Even Rosenblatt, who self-identified as in opposition to AI, secured continued placements in the 1950s and 1960s through grants via the Institute for Defense Analysis and the Office of Naval Research.³⁹

These affiliations galvanized the emerging research community. At the Dartmouth meeting, Herbert Simon, Allen Newell and J. Clifford Shaw introduced the results of their Logic Theory Machine, a virtual machine they designed at RAND in 1955–6 to prove thirty-eight theorems from *Principia Mathematica*. They positioned their results as evidence of having modelled human problem solving using an electronic digital computer. Ray Solomonoff, also in attendance, described the RAND machine as having solved 'the demo to sponsor problem', meaning the felt need to oblige funders with a proof of concept.⁴⁰ Minsky, too, noted that their 'inspiring progress' had had 'considerable effect on the direction of our work'.⁴¹ McCarthy called the RAND trio 'the stars of the show'.⁴² The Logic Theory Machine has since been designated, albeit contentiously, 'the first AI program' by existing histories of the field.⁴³

This group have collectively been described as the 'founding fathers' of artificial intelligence. As in historical accounts of colonialism in the United States, progenitor narratives describe early developments in terms of personalities and ideals, not access to resources.

37 Nilsson, op. cit. (34), p. 4.

38 John McCarthy, 'Dartmouth use of the IBM 704 to be located at MIT', n.d., 1, SC0524 ACCN 2013-247 Box 1; Miscellaneous from Notebook, Department of Special Collections and University Archives, Stanford University Libraries.

39 Penn, op. cit. (2), pp. 82–116.

40 Raymond J. Solomonoff, 'Untitled notes re: Wendy Conquest', 2005, Dartmouth AI Archives; Box A, Personal Archives of Grace Solomonoff, at <http://raysolomonoff.com/dartmouth/boxa/raywendypoints.pdf>; Grace Solomonoff, 'Ray Solomonoff and the Dartmouth Summer Research Project in Artificial Intelligence, 1956', Oxbridge Research, p. 19, Dartmouth AI Archives, Personal Archives of Grace Solomonoff (accessed 1 December 2016).

41 Marvin Minsky, 'A framework for artificial intelligence', 4 July 1956, Box A, Papers of Ray Solomonoff.

42 John McCarthy, 'The Dartmouth workshop – as planned and as it happened', John McCarthy's home page, at www-formal.stanford.edu/jmc/slides/dartmouth/dartmouth/node1.html (accessed 1 November 2019).

43 Boden notes that competing programmes in the field developed on both sides of the Atlantic. For a survey of instances see Margaret A. Boden, *Mind as Machine: A History of Cognitive Science*, 2 vols., Oxford: Clarendon Press, 2006, vol. 2, p. 705.

The term ‘founding fathers’ was likely introduced into AI’s historiography in 1979 by McCorduck’s *Machines Who Think*, which ‘forged the template for subsequent histories’, according to Mirowski.⁴⁴ McCorduck treated structural contingencies as extraneous, as Cohen-Cole has critiqued Margaret Boden for doing in her oft-cited *Mind as Machine*.⁴⁵ In the mid-1950s, RAND possessed perhaps the largest computing facility in the world. This and other corporate affordances to proto-AI, however, are treated as incidental, not constitutive. Neither IBM nor RAND figure in the list of AI’s ‘founding fathers’. Yet in contrast to cybernetics, institutional support for key researchers in what became AI was the rule, not the exception, during its rapid acceleration into an academic discipline.

In a 1957 lecture to the Operations Research Society of America, Simon and Newell self-reported what can be understood as an exception to this historiographical framing. They positioned their prototype as a milestone in the entangled histories of capitalism and manufacturing, stating,

For an appropriate patron saint for our profession, we can most appropriately look back a full half century before Taylor to the remarkable figure of Charles Babbage ... He was one of the strongest mathematicians of his generation, but he devoted his career to the improvement of manufacturing arts, and – most remarkable of all – to the invention of the digital computer in something very close to its modern form.⁴⁶

Their speech credited Adam Smith as the true inventor of the computer. While Smith’s theories on capitalism were not computational in practice, they argued, they were computational in principle, as Lorraine Daston corroborates.⁴⁷ Simon explained that thinkers like Gaspard de Prony and Charles Babbage had simply translated Smith’s ideas into hardware through iterative stages of development.

Simon and Newell’s institutional basis reflected this connection between technological engineering and social engineering. The RAND Corporation’s Systems Research Laboratory was among the world’s first experimental laboratories in management science. Further, as I have shown elsewhere, Simon, Newell and J. Clifford Shaw borrowed novel conceptual tools from precedents in administrative theory and formal logic to develop a functional machine language for the simulation of problem solving in a computer.⁴⁸ They repurposed elements of Simon’s influential 1947 text *Administrative Behavior* such as decision premises and means–ends analysis to shrink the program’s exposure to the complexity of the real world. *Behavior* posited that all activity in an organization could be reduced to explainable decision-making processes so long as the organization had a coherent, unchanging motive, specifically the capitalist prerogative to maximize profit. In other words, the Logic Theory Machine was a model of human behaviour, not of biology. Further, it was a model of a certain

44 Pamela McCorduck, *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*, San Francisco: W.H. Freeman, 1979, p. xii; Philip Mirowski, ‘Book review: McCorduck’s *Machines Who Think* after twenty-five years – revisiting the origins of AI’, *AI Magazine*, 2003, p. 135. ‘Founding fathers’ is rehashed in Nils Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements*, Cambridge: Cambridge University Press, 2010, p. 80; Howard Gardner, *The Mind’s New Science: A History of the Cognitive Revolution*, New York: Basic Books, 1998, p. 30. It is also used in more than four thousand other books related to AI, according to a search on Google Books in 2017.

45 Jamie Cohen-Cole, ‘Review of *Mind as Machine: A History of Cognitive Science* by Margaret Boden’, *Isis* (December 2008) 99(4), pp. 811–12.

46 Herbert A. Simon and Allen Newell, ‘Heuristic problem solving: the next advance in operations research’, *Operations Research* (1958) 6(1), pp. 1–2.

47 Lorraine Daston, ‘Enlightenment calculations’, *Critical Inquiry* (1994) 21(1), pp. 182–202; Daston, ‘Calculation and the division of labor, 1750–1950’, *Bulletin of the German Historical Institute* (Spring 2018) 62, pp. 9–30.

48 Penn, op. cit. (2), pp. 44–81.

type of behaviour, namely human problem solving, as seen through the prism of early twentieth-century symbolic logic and post-war American administrative logic.

The naturalization of AI's conceptual origins

Internalist historiographies by McCorduck and Boden bury these connections under what I argue are the rescripings of *animus nullius*, which foreground AI's empirical basis in neural – rather than social – phenomena and, in a related move, position AI as an academic science more than an industrial art. Ironically, Simon and Newell had resisted the term 'artificial intelligence' in the 1950s on the ground that it had been their results at RAND, which they categorized as 'complex information processing' and 'the next advance in operations research', that had lent validation to the abstract theories of Minsky, McCarthy and others. By the late 1950s, however, for no clear reason other than the traction the term 'AI' had gained at MIT and elsewhere, Simon and Newell began to use it, 'heuristic programming' and 'complex information processing' interchangeably. They, too, succumbed to the tendency to equivocate about the boundaries between metaphorical and literal terminology, between intelligence and problem solving, and between brains and computers.

These formative decisions mattered. In the United States, machine intelligence gained a new civility and sociability outside academia, spreading from research centres to newspapers, television and magazines. The vocabulary that early researchers chose to describe new techniques informed Americans' still plastic understandings of what was possible, and indeed desirable, in the emerging information age. Even if the earliest forms of 'AI' were definitively premised on a search for intelligence, which was inconsistently the case, it was also enmeshed in studies of search, storage and problem solving, as well as learning, vision, language, logistics and operations research, ways of knowing far removed from the sublimity of 'intelligence' and the brain.

By the late 1950s, allusions to neural fidelity were less concerned with any methodological commitment to close study of the brain than with a vaguely mentalistic way of positioning the field. Researchers were no longer sheepish about the need to qualify neural language. By the 1960s, 'automata' and 'brain models' became 'the artificial intelligence problem' and then just 'artificial intelligence'. I speculate that three broad trends prolonged ambivalence in AI about the exact meaning of suggestive neural metaphors. The first concerns patronage norms. In 1982, James Fleck attributed McCarthy, Minsky, Simon and Newell's sustained influence over the AI 'establishment' to their monopolized access to US military funding after 1963.⁴⁹ Fleck argued that this concentration, and the access to rarefied computational tools it enabled, was central to their mythos and to that of the Dartmouth meeting.⁵⁰ 'Until the early 1970s essentially all of DARPA's AI grants were given to the Massachusetts Institute of Technology and Stanford University'. Yarden Katz adds that 'between 1970 and 1980 MIT, Stanford, and Stanford Research Institute (SRI) received over 70 percent of the agency's AI funds'.⁵¹

Further research is needed to adequately disentangle the influence of commercial versus military funding in AI in the United States after 1960.⁵² Regarding its intricate

49 Fleck, op. cit. (16), p. 180.

50 On how the perceptrons controversy shaped this list see Jon Guice, 'Controversy and the state: Lord ARPA and intelligent computing', *Social Studies of Science* (1998) 28(1), pp. 103–38, 106; M. Olazaran, 'A sociological study of the official history of the perceptrons controversy', *Social Studies of Science* (1 August 1996) 26(3), pp. 611–59.

51 Yarden Katz, *Artificial Whiteness: Politics and Ideology in Artificial Intelligence*, New York: Columbia University Press, 2020, p. 25.

52 Philip Shiman and Alex Roland, *Strategic Computing: DARPA and the Quest for Machine Intelligence, 1983–1993*, Cambridge, MA: MIT Press, 2002; Katz, op. cit. (51), p. 9; Edwards op. cit. (31).

overlaps, consider that during the 1950s IBM earned more from its military contracts on the Semi-Automatic Ground Environment (SAGE) computing system than it did leasing out its own machines, and that IBM simultaneously contracted RAND's System Development Corporation (the lab where Newell, Simon and Shaw had met) to develop SAGE's software.⁵³ Colin Garvey's history of Japan's 1982 Fifth Generation Computer Systems project provides a telling point of contrast to evolving norms in the United States. Fifth Generation was 'the first national, large-scale AI R&D project to be free from military influence and corporate profit motives ... [realizing], in many ways, the norms of science often proclaimed – but rarely lived up to – by Western democracies'.⁵⁴ In response, the US government launched its billion-dollar Strategic Computing Initiative (SCI), which ran from 1983 to 1993. By 1985, it had allocated 50 million dollars each to public research laboratories and to industry.⁵⁵ Shiman and Roland's history of SCI as a search for 'machine intelligence' clarifies that the US government's aim was to 'remain competitive in the production of critical defense technologies'.⁵⁶ By this view, commercial and military aspirations for AI could, by the 1980s, be taken as two parts of the same whole.

A second trend, congruent with the first, was the rapid adoption of computing tools over the second half of the twentieth century. I speculate that the practicalities involved in this prodigious spread sedimented key terminology about computing that might otherwise have remained pliant to scholarly critique and refinement.⁵⁷ While only two digital electronic computers were in operation in the United States in 1950, 243 machines were operational in 1955, with 5,400 in 1960, 25,000 in 1965 and 75,000 in 1970.⁵⁸ After 1960, mention of 'automatic programming' and 'automatic coding' techniques became discussions about 'software'.⁵⁹ How this broad spread influenced rhetoric about 'AI' has yet to receive sustained historical treatment.⁶⁰

A third trend prolonging ambivalence about suggestive neural metaphors concerns the sustaining effect of isomorphisms between methods in AI and mainstream economics. Emerging scholarship suggests that, at a high level, efforts at digitization followed efforts at financialization. Devin Kennedy, for instance, explores how foundational theory for how to sort and manage complexity in computer science mirrored approaches for how to sort and manage capital.⁶¹ Similarly, with regard to the history of software development after the 1970s, Laine Nooney argues, 'Throughout it all, we can clearly see the intractable role financial speculation and the construction of markets played in people's desire to even imagine what shape innovation might take.'⁶² In relation to 'AI', Rosenblatt's perceptron, a foundational contribution to machine learning, explained the

53 Thomas Haigh and Paul E. Ceruzzi, *A New History of Modern Computing*, Cambridge, MA: MIT Press, 2021, p. 88.

54 Colin Garvey, 'Artificial intelligence and Japan's fifth generation: the information society, neoliberalism, and alternative modernities', *Pacific Historical Review* (November 2019) 88(4), pp. 619–58, 621.

55 McCorduck, op. cit. (44), p. 429; also explored in Shiman and Roland, op. cit. (52), p. 195.

56 Shiman and Roland, op. cit. (52), p. 5.

57 For a survey of AI's critics see Shunryu Colin Garvey, 'Unsavoury medicine for technological civilization: introducing "Artificial Intelligence & its discontents"', *Interdisciplinary Science Reviews* (3 April 2021) 46(1–2), pp. 1–18.

58 Nathan Ensmenger, *The Computer Boys Take Over: Computers, Programmers, and the Politics*, Cambridge, MA: MIT Press, 2010, p. 28.

59 Martin Campbell-Kelly and Daniel D. Garcia-Swartz, 'Pragmatism, not ideology: historical perspectives on IBM's adoption of open-source software', *Information Economics and Policy* (August 2009) 21(3), pp. 229–44, 234.

60 Towards this end, I have contributed: Penn, op. cit. (2), pp. 116–58.

61 Devin Kennedy, 'Virtual capital: computers and the making of modern finance, 1929–1975', PhD dissertation, Harvard University, 2019.

62 Laine Nooney, *The Apple II Age: How the Computer Became Personal*, Chicago: The University of Chicago Press, 2023, p. 261.

irreducibly complex phenomena of the mind using Hayek's notion of decentralized market structures. Tracing the mathematics of 'minds' and 'markets' forward, one finds similar overlaps in the history of the loss function, in the Black–Scholes options pricing model and in network science generally.⁶³

That the search for a theory of machine intelligence derived from the empirics of social phenomena, and not from neurophysiology alone, became stubbornly apparent as research vogues shifted after the 1960s. In 1976, Newell and Simon posited, perhaps at the peak of symbolic AI research, that 'a physical symbol system has the necessary and sufficient means for general intelligent action'.⁶⁴ Unconvinced by the feasibility of this goal, those in the emerging, alternative tradition of 'expert systems' in the 1960s to 1980s set out to encode human knowledge, not human reasoning, using databases of 'if-then' rules. Edward Feigenbaum, who coined the term 'expert systems', wrote, 'The problem-solving power exhibited by an intelligent agent's performance is primarily the consequence of its knowledge base, and only secondarily a consequence of the inference method employed. Expert systems must be knowledge-rich even if they are methods-poor ... The power resides in the knowledge.'⁶⁵

In the 1970s to 1990s, expert systems found both novel commercial uses and piercing critique. Robert Cooper, director of DARPA, took pains to 'transition' expert-systems research from the lab to the marketplace and to connect university researchers to industry.⁶⁶ Critique followed. In the 1990s, the anthropologist Diana E. Forsythe challenged researchers' self-reported abilities to extract, formalize and model so-called 'domain-specific' knowledge using interviews with human experts.⁶⁷ 'Whereas anthropologists devote considerable energy to pondering methodological, ethical and philosophical aspects of field research', she wrote, 'knowledge acquisition seemed to be undertaken in a rather unexamined way. Asked how they went about the task of gathering knowledge for their expert systems, the knowledge engineers I met tended to look surprised and say, "We just do it."⁶⁸

While expert systems fell out of favour in AI research in the 1990s, the positivist assumption that rendered knowledge a self-evident, structured and accessible 'thing' that could be 'extracted' carried into the next vogue of AI – machine learning – through technical vocabulary about 'domains'.⁶⁹ This term finds ubiquitous use in contemporary machine learning. Ribes *et al.* position its rhetorical power as a claim to 'identify, demarcate and characterize spheres of worldly action or knowledge' that in daily use implies, necessarily, that 'there is a more general, even universal, method or technique ... [that

63 Castelle, 'Are neural networks neoclassical? Utility, loss, and cost from Wald to Tensor', at www.youtube.com/watch?v=IKBpOQ3UdGw (accessed 28 October 2023). Thanks to Yaqub Chaudhary for Orit Halpern, 'Financializing intelligence: on the integration of machines and markets', *E-Flux Architecture*, March 2023, at www.e-flux.com/architecture/on-models/519993/financializing-intelligence-on-the-integration-of-machines-and-markets; Wendy Hui Kyong Chun and Alex Barnett, *Discriminating Data: Correlation, Neighborhoods, and the New Politics of Recognition*, Cambridge, MA: MIT Press, 2021, pp. 70–114.

64 Allen Newell and Herbert A. Simon, 'Computer science as empirical inquiry: symbols and search', *Communications of the ACM* (March 1976) 19(3), pp. 113–26, 116.

65 Edward A. Feigenbaum, 'Knowledge engineering: the applied side of artificial intelligence', *Annals of the New York Academy of Sciences* (November 1984) 426(1), pp. 91–107, 101.

66 Shiman and Roland, *op. cit.* (52), p. 4.

67 On the rigidity of the conduct required to use expert systems see Stephanie A. Dick, 'Coded conduct: making MACSYMA users and the automation of mathematics', *BJHS Themes* (2020) 5, pp. 205–24.

68 Diana E. Forsythe, 'Engineering knowledge: the construction of knowledge in artificial intelligence', *Social Studies of Science* (August 1993) 23(3), pp. 445–77, 447.

69 On machine learning pre-dating AI see Aaron Plasek, 'On the cruelty of really writing a history of machine learning', *IEEE Annals of the History of Computing* (October 2016) 38(4), pp. 6–8.

can be used] across many, and sometimes all, domains'.⁷⁰ In his study of data positivism since the Second World War, Matthew Jones points to Rosenblatt's perceptron theory as epitomizing a branch of instrumentalist statistics that sought technique-as-explanation, meaning functions that fit the data rather than functions that fit a corresponding law of nature.⁷¹ This instrumentalist view of knowledge distanced machine-learning researchers from the empirics of the material world, even as they draped their techniques in the suggestive language of neural mechanisms, as the name 'perceptron' indicates.

In the 2010s, American technology conglomerates mobilized to fuel widespread media speculation of a looming 'AI revolution' in which Yarden Katz has argued that AI doubled for 'a confused mix of terms – such as "big data," "machine learning," or "deep learning" – whose common denominator is the use of expensive computing power to analyse massive centralized data'.⁷² Behind this marketing, Gürses argues, was a concerted push by the global finance industry to position American technology conglomerates as a safe harbour for large-scale institutional investment following the collapse of real-estate investments during the 2007–8 financial crisis. To satisfy shareholders' expectation of a return, the largest cloud service providers – Amazon, Google, Microsoft – turned to the 'organizational capture' of hospital, university and government tech budgets, among other small and medium-sized businesses.⁷³ As with Facebook's efforts to monopolize the marketplace for global advertising budgets, American cloud providers pressed to capture organizations' technology infrastructure budgets by enlisting them into pay-as-you-go hardware and software services accessed via distributed data centres known as 'the cloud'. Per Katz, rhetoric about 'AI' figured centrally here as a euphemism for private cloud infrastructure.⁷⁴

To incentivize and prolong adoption of their proprietary infrastructure, tech conglomerates invested heavily in the development of machine-learning techniques that could, in operation, enclose regions of common knowledge online. This time, they indulged AI's sociotechnical imaginary to veil acts of seizure as acts of novel transformation or discovery.⁷⁵ Blackwell calls this 'institutionalized plagiarism'.⁷⁶

Fundamental to such seizures is a lack of case law, which compounds existing confusion over the provenance of authorship, the role of human contributors in creating input data, and the degree to which generative AI outputs are original and transformative. A class action lawsuit filed against OpenAI in July 2023 rejects the idea that their ChatGPT model is capable of genuine novelty. 'A large language model's output is ... entirely and uniquely reliant on the material in its training dataset', they reason.⁷⁷ According to

70 David Ribes, Andrew S. Hoffman, Steven C. Slota and Geoffrey C. Bowker, 'The logic of domains', *Social Studies of Science* (June 2019) 49(3), pp. 281–309, 282.

71 Matthew L. Jones, 'How we became instrumentalists (again): data positivism since World War II', *Historical Studies in the Natural Sciences* (November 2018) 48(5), pp. 673–84, 678.

72 Yarden Katz, 'Manufacturing an artificial intelligence revolution', *SSRN Electronic Journal*, 2017, at <https://ssrn.com/abstract=3078224>, p. 2; for crossover between machine intelligence and the economics of surveillance see Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for the Future at the New Frontier of Power*, London: Profile Books, 2019.

73 Seda Gürses, 'How Big Tech captured our public health system', interview by Arun Kundnani, 18 May 2022, at www.youtube.com/watch?v=RBA1bZ2fKKc&t=1581s (accessed 1 June 2022).

74 Left open is whether this marketing qualifies as agnotology. See Naomi Oreskes and Erik M. Conway, *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*, New York: Bloomsbury Press, 2010.

75 Sheila Jasanoff and Sang-Hyun Kim (eds.), *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*, London: The University of Chicago Press, 2015.

76 Alan Blackwell, *Moral Codes: Designing Software without Surrender to AI*, Cambridge, MA: MIT Press, 2022, p. 11.

77 Sarah Silverman, Christopher Golden and Richard Kadrey, *Silverman v. OpenAI, Inc.*, 3:23-cv-03416, District Court, N.D. California, 7 July 2023, ECF No. 1, No. 3:23-cv-03416, at www.courtlistener.com/docket/67569254/silverman-v-openai-inc (accessed 18 September 2023).

this view, OpenAI has made illegal use of copywritten material. The company's response, in August 2023, situates the case around the science of inference. 'As Plaintiffs acknowledge, the ability of digital software to [interpret user prompts and generate convincingly naturalistic text outputs] reflects a legitimate scientific advancement in the field of artificial intelligence, which seeks to "simulate human reasoning and inference."'78

This lawsuit captures aspects of the evolving character of *animo nullius* after the 2010s. Unlike in the mid-century, it is now corporations, alongside some researchers, who stand to benefit materially from equivocation over the perceived authority and intellectual genealogy that underwrites AI. OpenAI's valuation fluctuates in line with this perception; at the time of going to press it is the most valuable start-up in America. Post-war IBM sought only outputs from AI, not its brand. Another difference is that debate over these scientific merits has shifted from labs and markets to courtrooms, too. That territories seized under *terra nullius* remain stolen today, centuries later, bears out how misinterpretations of the past can calcify into law and prolonged settlement. Discerning publics – including judges, legislators and dispossessed authors – would be forgiven for assuming, naively, that the statistical 'intelligence' and 'learning' on display in post GPT-era systems derive from, say, the empirics of Darwin rather than, in my account, the contingent social order elaborated by Smith. What else would history suggest?

Birch, who theorizes the rise of assetization in modern technoscience, characterizes this form of economic system as rentiership. Unlike with the valuation of a commodity, which is generally extinguished upon its consumption, ownership and control of an asset extend one's rights to copies and derivatives, which adds scope for control and extraction. 'No one can tell you how to wear a coat you have bought, but they can tell you how to use the copyrighted music or data they have sold you', he writes.⁷⁹ *Animo nullius* figures in rentiership economics because the value of an asset requires 'active and ongoing organization, governance, and management ... it is not a passive process'.⁸⁰ In my account, *animo nullius* animates AI's value as an asset by granting those who invoke it an outwardly credible and durable (that is, historical) claim to the authority of a prestige metaphor.

Data colonialism

In my introduction, I cited Mejias and Couldry's interest in how these emerging branches of capitalism intersect with data colonialism. I will end by offering three conjectural parallels that emerge from my account. The first is an observable convergence between state and corporate interests over privatizing the costs of speculation (and initial seizures) amidst imperial expansion. State involvement is central to the durability of the extractive project outlined above because, per Birch, 'assets are legal constructs, in that ownership and control rest on the state enforcement of property and control rights'.⁸¹ In *Empire of Cotton*, Sven Beckert argues that the commodification, processing and global trade of cotton illuminate the mechanics of historical colonialism and the Atlantic slave trade, wherein empires backed merchants and colonial corporations with military force to facilitate

78 OpenAI L.L.C., *Silverman v. OpenAI, Inc.*, 3:23-cv-03416, District Court, N.D. California, 28 August 2023, ECF No. 32, Motion to Dismiss – Document #32 15 (n.d.), at www.courtlistener.com/docket/67569254/silverman-v-openai-inc (accessed 18 September 2023).

79 Kean Birch (ed.), *Assetization: Turning Things into Assets in Technoscientific Capitalism*, London: MIT Press, 2020, p. 12.

80 Kean Birch, 'Technoscience rent: toward a theory of rentiership for technoscientific capitalism', *Science, Technology, & Human Values* (January 2020) 45(1), pp. 3–33, 18.

81 Birch, op. cit. (80), p. 12.

expansion.⁸² I suggest that the race to enclose the knowledge commons and assetize key analytical processes to the benefit of the global rentier class (and their respective nation states) helps to structure how data colonialism will shape capitalism in the era ahead.

A second parallel, which puts tension on the first, is states' backing of corporate self-governance and claims to the commons. State-backed corporate agency was decisive to the settler colonial project in 'New England', the land upon which, centuries later, AI would initially be developed at MIT. In that instance, the mechanics of self-governance in colonial corporations helped to prototype and, upon maturation, to stabilize and cohere America's legislative independence from the British Empire.⁸³ Possession via occupation succeeded wildly for the Puritans of New England, and for the Massachusetts Bay Company.⁸⁴ Suggestive of a disanalogy with *terra nullius*, however, Birch notes that it is contract law rather than property rights that is central to the ownership of intangible assets such as machine-learning models.⁸⁵ In any case, corporate self-governance is sure to remain an area of contestation as notions of sovereignty and power intersect with the global use of proprietary inference models across governments, universities, hospitals and other civic institutions.

One final line of continuity concerns modes of capture. Analysing the relationship between *terra nullius* and historical colonialism, Buchan centres the role of subversive acts of 'trafficking', meaning the provision of gifts, favours and symbols of office that – as tokens of exchange – subtly coerced consent from Indigenous peoples that European colonists encountered to initiate trade relations. Over decades and centuries, these disingenuous relations reduced to modes of assimilation, subordination and then subjugation. While capture is already a concern in areas of AI critique, its relation to historiography and coloniality deserves particular attention.⁸⁶ Indigenous scholars have shown how historiographical claims naturalize dispossession. Jean M. O'Brien's *Firsting and Lasting: Writing Indians out of Existence in New England*, for instance, charts how colonizers narrativized Indigenous peoples as extinct in order to assert and mythologize their own sovereignty, ancestry and social order. Notably, Zoe Todd advances a similar critique of the history of science itself.⁸⁷

Conclusion

As with *terra nullius*, *animus nullius* threatens the commons with enclosure and private management. Both merit rebuke. In this paper, I have resituated historiographies of AI in the history of corporate colonial expansion to account for the fields' imbrications in systems of power that prefigure its emergence in the mid-twentieth century. I have shown how, as far back as the 1950s, corporate efforts to incorporate digital media into mainstream

⁸² Sven Beckert, *Empire of Cotton: A Global History*, New York: Vintage Books, 2015.

⁸³ Herbert L. Osgood, 'The corporation as a form of colonial government. I', *Political Science Quarterly* (1896) 11 (2), pp. 259–77; Osgood, 'The corporation as a form of colonial government. II', *Political Science Quarterly* (1896) 11 (3), pp. 502–33.

⁸⁴ James Muldoon, 'Colonial charters: possessory or regulatory?', *Law and History Review* (May 2018) 36(2), pp. 355–81.

⁸⁵ Birch, op. cit. (79), p. 25.

⁸⁶ Bruce Buchan, 'Traffick of empire: trade, treaty and terra nullius in Australia and North America, 1750–1800', *History Compass* (March 2007) 5(2), pp. 386–405. Rodrigo Ochigame, 'The invention of "ethical AI": how big tech manipulates academia to avoid regulation', *The Intercept*, 20 December 2019, <https://theintercept.com/2019/12/20/mit-ethical-ai-artificial-intelligence>; Meredith Whittaker, 'The steep cost of capture', *Interactions* (November 2021) 28(6), pp. 50–5.

⁸⁷ Jean Maria O'Brien, *Firsting and Lasting: Writing Indians out of Existence in New England*, Minneapolis: University of Minnesota Press, 2010. Zoe Todd, 'An Indigenous feminist's take on the ontological turn: "ontology" is just another word for colonialism', *Journal of Historical Sociology* (March 2016) 29(1), pp. 4–22.

knowledge practices (and, in doing so, to sell standardized computing infrastructures, from IBM to Microsoft) have intertwined with positivist interpretations of mind-as-computer metaphors as both literal and perpetually looming. Even when these overlaps were the products of opportunism by one side or the other, privatization appears to be the outcome. This speaks to the political economy within which AI was formed – significant elements of which are usually neglected by practitioners, historians and media commentary, however critical, on AI. For historians, the decision to characterize ‘AI’ as primarily a science of the mind, rather than a science of operations and administration, is, in my view, just that – a decision, and one that neglects the breadth of AI’s origins and helps blind us to its imbrications.

Acknowledgements. Special thanks to Richard Staley, Stephanie Dick, Mustafa Ali, Sarah Dillon, Helen Curry, Jon Agar and Simon Schaffer for their feedback on previous versions of this paper, and to Matthew Jones for similar guidance. Thank you also to Susie Gates and the many who brought such life to our Histories of AI Sawyer Seminar.