‘INTERDISCIPLINARY INTERACTIONS’: NIPS 1987 AND THE POLITICS OF NEURAL NETWORKS

UNIVERSITY OF CAMBRIDGE

DEPARTMENT OF HISTORY AND PHILOSOPHY OF SCIENCE

Harry Law

Table of contents

Introduction: 'Interdisciplinary interactions'

On 8 November 1987 the Sheraton Tech Center Hotel in Denver, Colorado opened for the inaugural Conference on Neural Information Processing Systems – Natural and Synthetic (NIPS). Six hundred 'neural network' researchers working in neurobiology, physics and information processing gathered to mark the beginning of a conference designed, according to its programme, to promote 'interdisciplinary interactions'.[1] 'The conference was a huge hit, and everyone realized that we have something valuable,' wrote NIPS programme chairman Yaser Abu-Mostafa in reflections on the event.[2] More than perhaps a passing remark on the conference, Abu-Mostafa's comments raise an important question of the Denver meeting: in what way can 'everyone' be understood?

This paper explores the role of interdisciplinarity in the remaking of behaviourism by the connectionist branch of artificial intelligence (AI). I examine the 1987 Conference on Neural Information Processing Systems or 'NIPS', the foundational event in what became an influential machine learning conference series but about which little historical literature has been written.[3] My paper draws on science and technology studies researcher Alison Adam's 'view from nowhere' critique, which argues that AI practitioners have consistently assumed to 'speak for some universal yet never articulated subject, nowhere yet everywhere at the same time.'[4] Applying the work of Thomas Nagel to AI, Adam argues that researchers have repeatedly failed to consider the identity of the 'thinking' subject, problematically positioning AI systems as representative of a neutral, universal perspective.[5] Though many critical approaches have been applied to AI such as standpoint theory and post-colonial theory, I consider the view from nowhere to explore the relationship between universalism and interdisciplinarity at NIPS.[6] I examine conference proceedings, press interviews, and administrative materials to argue that the 'interdisciplinary interactions' championed at NIPS reflected a remaking of behaviourism by the connectionist community in the United States in the late 1980s.

My paper is organised as follows. After examining the contested character of notions of cognition, computing, and society I provide a brief description of the history of behaviourism, connectionism, and the launch of NIPS in

[1] Anon., Abstract of Papers. IEEE Conf. on "Neural Information Processing Systems - Natural and Synthetic." Institute of Electrical and Electronics Engineers, June 30, 1989, 2. https://apps.dtic.mil/dtic/tr/fulltext/u2/a212697.pdf. Anon., "Advance Program." Institute of Electrical and Electronics Engineers, November 1987.
[2] Yaser Abu-Mostafa, "The first NIPS (now NeurIPS)," last modified 19 May 2020. http://work.caltech.edu/nips.
[3] Margaret Boden, *Mind as Machine,* (Oxford: Clarendon Press, 2006) 960 and 961.
[4] Alison Adam, *Artificial knowing: Gender and the Thinking Machine* (London: Routledge, 2006) 4.
[5] Thomas Nagel, *The View from Nowhere*, (Oxford: Oxford University Press, 1986) 3.
[6] Stephen Cave and Kanta Dihal, "The Whiteness of AI," *Philosophy and Technology*. 33, 685–703 (2020). Francesca Ferrando, "Is the post-human a post-woman? Cyborgs, robots, artificial intelligence and the futures of gender: a case study," *European Journal of Futures Research.* 43 (2014) 4.

Denver in 1987. I apply, and build-on, Adam's view from nowhere argument to surface three ways in which NIPS represented a remaking of behaviourism within the connectionist school of AI through rhetorical and technical appeals to the universality of research. First, notions of interdisciplinarity embodied by NIPS legitimised behaviourist theories of artificial neural networks and brain function challenged by experts in other domains. Second, research presented at the conference reflected the way in which neural networks encoded behaviours that privileged perspectives of masculinity, whiteness, and affluence by, for example, making general claims about cognition through the extrapolation of experimental results using a single male test subject. Third, the use of the language of neural networks by those associated with NIPS to describe systems of social organisation sharpened connectionist social logics that echoed the utopian character of Skinnerian behaviourism.

One: From Behaviourism to Connectionism

Sheila Jasanoff has demonstrated the way in which collectively held futures, expressed as 'sociotechnical imaginaries', are used to describe what the future 'could' look like and prescribe what the future 'should' look like.[7] As such, futures—including those of cognition, computing, and society in the United States in the twentieth century—can be understood as generated through an 'unstable field of language, practice and materiality' in which agents compete to represent progress and deterioration.[8] In 1961, for example, American television network CBS developed 'The Thinking Machine' in which MIT computer scientist Jerome Wiesner made strong claims about the relationship between the future of computing and the future of American society, noting that machine intelligence will 'spark a revolution that will change the face of the earth'. [9] Wiesner was not alone. A 1958 *Mechanix Illustrated* magazine article described an IBM computer as a 'giant brain', while Marvin L. Minsky, an influential American AI researcher, also at MIT, speculated in a 1966 *Scientific American* article that the existence of sophisticated machine intelligence would render human activities and aspirations 'changed utterly'.[10] I draw on these examples to underscore the contested, interconnected nature of the rhetoric of mind and machine in the United States after the mid-twentieth century, which, as this paper will demonstrate, intersected with force once again at NIPS in 1987.

---

[7] Sheila Jasanoff and Sang-Hyun Kim, "Containing the Atom: Sociotechnical Imaginaries and Nuclear Power in the United States and South Korea," *Minerva* 47, no. 2 (June 2009): 139.
[8] Nik Brown and Brian Rappert, *Contested Futures: A Sociology of Prospective Techno-Science* (Oxford: Routledge, 2017) 5.
[9] CBS Tomorrow, "The Thinking Machine," broadcast 1961, MIT Museum Collections, 09:30. https://techtv.mit.edu/videos/10268-the-thinking-machine-1961---mit-centennial-film.
[10] Marvin Minksy, "Artificial Intelligence," *Scientific American*, September 1966, 246. Robert Strother "Thinking machines are getting smarter," *Mechanix Illustrated*, 10.

Though fundamental to the promotion of machine intelligence, imaginaries surrounding the brain and cognition had also long been used to describe possible systems of social organisation. Consider the American psychologist B.F. Skinner, whose influential theory of behaviourism posited that human will and agency could be understood as products of external stimuli. In 1948, as Skinner studied operant conditioning—in which the strength of an individual's behaviour is modified by reinforcement and controlled by environmental stimuli—he published his utopian novel *Walden Two* that examined the ways in which entire societies could be engineered based on the control of behaviour.[11] Idealisations of cognitive functioning, as the book makes clear, can be used to describe how societies might function and prescribe how societies ought to function. This essay, through a close study of NIPS 1987, shows that interconnected imaginaries of cognition, computing, and society contained shared conceptions of knowledge that underpinned the remaking of behaviourism within the connectionist cosmology.

A brief consideration of the historiographical context is needed to situate NIPS 1987 within the genealogy of AI research. Popular histories of AI, some written by practitioners themselves, have tended to cluster research on 'thinking machines' into one of two technical schools. First, the so-called 'symbolic' approach in which systems are built using hard-coded rules based on the manipulation of expressions. Second, the 'connectionist' perspective, which argues that systems ought to mirror the interaction between neurons of the brain. Idealised 'artificial neural networks' or 'neural networks' belong to this group. Scholars like Lucy Suchman, Stephanie Dick, Matthew L. Jones, Aaron Plasek, and Jonnie Penn have stressed the need to consider these accounts in relation to the material, political, and social circumstances surrounding computation and data.[12] This paper follows in that tradition in its examination of the connectionist school, which was generally recognised to have suffered a dearth in interest during the 1950-60s prior to a resurgence in the mid-1980s.[13] MIT researchers Michel Rappa and Koenraad Debackere found that the rate of growth of the connectionist community increased by 60 per cent a year during the 1980s as membership expanded five-fold, while 75 per cent of neural network researchers polled in 1990 entered the field between 1984 and 1990.[14] No fewer than four major events dedicated to the discussion of neural

[11] B. F. Skinner, *Walden Two*. (New York: Macmillan, 1976) xvi.
[12] Aaron. Plasek, "On the Cruelty of Really Writing a History of Machine Learning," *IEEE Annals of the History of Computing,* 38 (2016): 6-8. Stephanie Dick, 'Of Models and Machines: Implementing Bounded Rationality,' *Isis* 106, no. 3 (September 2015): 623–34 Matthew L. Jones, *Reckoning with Matter: Calculating Machines, Innovation, and Thinking about Thinking from Pascal to Babbage*, (Chicago: The University of Chicago Press, 2016). Lucille Alice Suchman, *Human-Machine Reconfigurations: Plans and Situated Actions*. (Cambridge: Cambridge University Press, 2007).
[13] Interest in connectionist approaches is generally considered to have waned following the publication of the 1969 book *Perceptrons: An Introduction to Computational Geometry* by Minsky and Papert. Mikel Olazaran, "A Sociological Study of the Official History of the Perceptrons Controversy," *Social Studies of Science* 26, no. 3 (August 1996): 611-659.
[14] Michel Rappa and Koenraad Debackere, *The Emergence of a New Technology: The Case of Neural Networks*, Alfred P. Sloan School of Management, 1989, 7. Michel Rappa and Koenraad Debackere, "International survey on the neural network research community," Alfred P. Sloan School of Management, 11.

network research were held in 1987, with each of the conferences hosting an inaugural edition in the mid-1980s.[15] The growth in conferences catalysed the sort of 'interdisciplinary interactions' I argue contributed to the newfound vitality of connectionism and, with it, a remaking of behaviourism. The launch of NIPS in 1987 provides a window into this transformative episode.

Technical advances, sociological factors and a groundswell of industrial and military interests played a role in returning connectionism to prominence in the 1980s. In a paper presented in Denver, University of Calgary computer scientist Bruce MacDonald argued that the resurgence was driven by downward pressure from dissatisfaction with the symbolic manipulation paradigms of the 1960s and 1970s and upward pressure from the emergence of computers with significant processing power.[16] In a 1988 paper, MIT computer scientist Seymour Papert suggested that the 'behaviorist process of external association of stimuli with reinforcements' generated a 'cultural resonance' between behaviourist interpretations of mind and the connectionist credo.[17] The 'evidence' provided by neural networks for behaviourist understandings of mind proved to be a self-reinforcing phenomenon; neural networks 'confirmed' the behaviourist perspective, which, in turn, drove interest in connectionist models that benefitted from suggestions that they accurately represented human cognition. Meanwhile, sponsorship of NIPS by U.S. WEST Advanced Technologies, the Office of Naval Research, and the Air Force Office of Scientific Research ensured defence organisations enjoyed oversight of conference literature.[18] U.S. defence agency DARPA noted a 'frantic' interest in connectionist systems in its 1987 Neural Network Study, while American military interests were associated with all four major neural network conferences in 1987.[19] Corporate backing for NIPS was provided by Bell Labs, with IBM and Hecht-Nielsen Neurocomputers providing additional support.[20] Amidst a backdrop of sociological and technological factors, the introduction of corporate and military patronage energised a connectionist rally in which NIPS emerged as a leading forum for the discussion of research in 1987.

---

[15] T.J. Schwartz, "1987, the Neural Year in Review," *Synapse Connection* (February), 1988 1 and 11-12. Mikel Olazaran, "Historical Sociology of Neural Network Research", PhD Thesis; University of Edinburgh, 1991, 283-268.
[16] Bruce MacDonald, "Connecting to the Past." In *Proceedings of the 1987 International Conference on Neural Information Processing Systems*, 505–514 ed. Dana Anderson (New York: American Institute of Physics, 1988) 1. Hereafter "Proceedings*"*.
[17] Seymour Papert, "One AI or Many?" In *The Artificial Intelligence Debate: False Starts, Real Foundations*, ed. Stephen Graubard (Cambridge: MIT Press, 1988) 159.
[18] Anon., "Advance Program."
[19] Mikel Olazaran, "Historical Sociology of Neural Network Research" (PhD Thesis, University of Edinburgh, 1991) 282-292.
[20] Anon., "Advance Program."

Two: Rhetoric of Mind and Machine

NIPS conference papers made rhetorical appeals to perceived equivalences between artificial neural networks and the biological realities of neuroscience. Researchers in Denver argued that neural networks, which 'learned' through the interactions of highly abstract 'artificial neurons' known as parameters, echoed the Skinnerian process of operant conditioning. John Denker and his Bell Labs colleagues recast the brain as an 'organic computer' in a paper presented at the conference, while computer scientist Dan Hammerstrom of the Oregon Graduate Center noted that artificial neural networks had started to 'exhibit a broad range of cognitive behavior' in separate NIPS research.[21] Enis Ersü, CEO of German computer vision specialist ISRA Systemtechnik GmbH, underscored the nature of the relationship in his conference paper (Figure 1.0) by comparing his LERNAS neural network with the cerebral cortex.[22] A further NIPS paper by University of California neuroscientist Gary Lynch explained that artificial environments provided a framework for understanding the biological properties of neurons, which would aid in the design of network architectures inspired by biology.[23] By Denker, Hammerstrom, Ersü, and Lynch's accounts, the reflexive nature of the relationship between neuroscience and artificial neural networks precipitated a process of mutual reinforcement; the former used simulations to circumvent the limits of anatomical study while the latter benefitted from association with the prestige and authority afforded by the mind. Though not every paper from the conference programme explicitly adopted this approach, none challenged the use of rhetoric in this way or offered alternative understandings of 'natural and synthetic' neural networks.

[21] John Denker et al., "Microelectronic Implementations of Connectionist Neural Networks," In "Proceedings," 515–523. Dan Hammerstrom. "The Connectivity Analysis of Simple Association," in "Proceedings," 338–347.
[22] Enis Ersü and Henning Tolle, "Hierarchical Learning Control - An Approach with Neuron-Like Associative Memories," In "Proceedings," 249-261.
[23] Gary Lynch et al., "Partitioning of Sensory Data by a Cortical Network," In "Proceedings," 317–337.

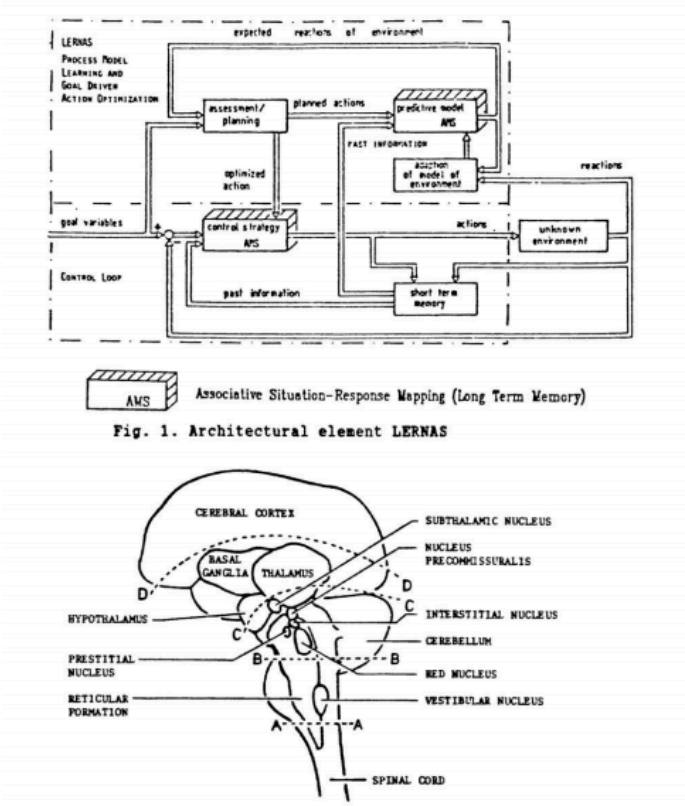

Figure 1.0. An excerpt from Enis Ersü's NIPS paper comparing the LERNAS neural network to the human brain.[24]

These assumptions did not stay within the walls of the Sheraton Tech Center. Speaking to the *Pittsburgh Post-Gazette* six months prior to NIPS in May 1987, connectionism pioneer David Rumelhart said he had been 'astounded by the similarities' between neural network research (depicted in Figure 2.0) and 'actual memory processing' in the human brain.[25] 'Science is stealing tricks from the brain,' quipped Lynch to the *Schenectady Gazette* in March of 1987. '[Human brains] do the magic they do from the cortex. We're taking just a small piece of that magic and building it into silicon.'[26] The *New York Times* reported in September 1987 that NIPS 'neurobiology liaison' James Bower, whose remit was to encourage the neurobiology community to engage with the conference, predicted that as neural networks became more sophisticated anatomists would develop theories of the brain according to evidence uncovered by connectionist researchers.[27] Imaginaries of computer and brain––communicated by Rumelhart, Bower, and Lynch and published by a sympathetic media—thus embodied a potent strain of universalism inherent within the connectionist dogma. *Every* mind was, in essence, a computer, lay

[24] Enis Ersü, "Hierarchical Learning Control," 258.
[25] Henry Peirce, "Meetings of the mind," *Pittsburgh Post-Gazette*, May 11, 1987 28.
[26] Delthia Ricks, "Computer Researchers Look to Human Brain." *Schenectady Gazette*, March 5, 1987 20.
[27] Andrew Pollack, "More Human Than Ever, Computer is Learning to Learn," *New York Times*, September 15, 1987.

readers were told. Stated differently, just as Adam's view from nowhere argues that a failure to consider the thinking subject enables the values and ideals of researchers to inadvertently contaminate neural network architecture, so, too, did a failure to critically analyse connectionist systems—stymied by mutually reinforcing conceptions of mind and machine—produce a technical ecology in which the seed of behaviourism might flourish.

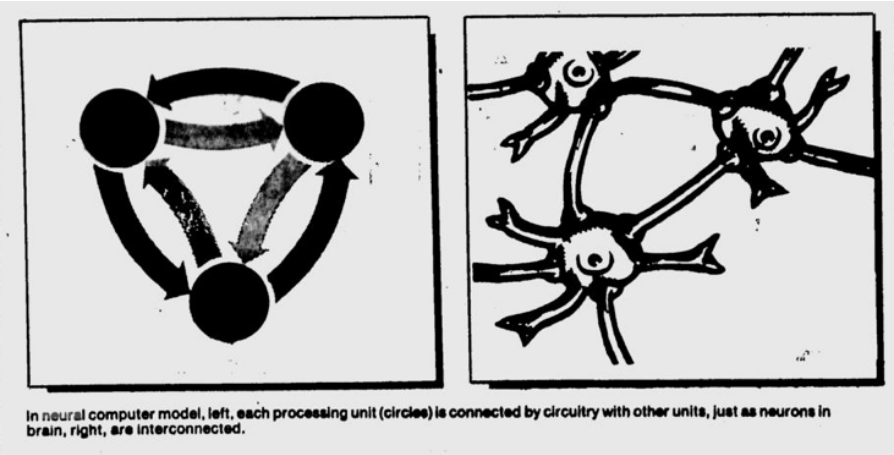

Figure 2.0. An excerpt from the Pittsburgh Post-Gazette *depicting highly abstract 'parameters' in a neural network as akin to the brain's biological neurons.*[28]

Not all were convinced. In 1988, a year after NIPS, Rutgers Center for Cognitive Science researchers Jerry Fodor and Zenon Pylyshyn published a comprehensive rebuke of the 'lure of connectionism'. Arguing that systematic relations amongst the brain's cognitive capacities are causally sensitive to the constituent structure of mental representations, Fodor and Pylyshyn put forward that the mind is not connectionist at the cognitive level.[29] Neural network researchers, including NIPS attendee and University of Colorado computer scientist Paul Smolensky, responded to the paper by arguing that the problem of systematicity could be solved by the existence of 'distributed' mental representations.[30] The paper provoked a response from Fodor and his colleague Brian McLaughlin that returned compositional structure was irrelevant to the systematicity problem.[31] I reference the debate to call attention to the omission of such critiques from the NIPS programme and to underscore the contested nature of connectionist understandings of cognition. The exclusion of research that challenged the core tenants of connectionism demonstrates the fragility of the 'evidence' that artificial neural networks provided for brain

[28] Henry Peirce, "Meetings of the mind," 28.
[29] Jerry Fodor and Zenon Pylyshyn, "Connectionism and cognitive architecture: A critical analysis," *Cognition*, 28 (1988): 3–71.
[30] Paul Smolensky, "The constituent structure of connectionist mental states: A reply to Fodor and Pylyshyn," *The Southern Journal of Philosophy*, 26 (1988):137–161.
[31] Jerry Fodor and Brian McLaughlin, "Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work," *Cognition*, 35 (1990):183–204.

function. The event's 'interdisciplinary interactions' were 'interdisciplinary' in a distinctly narrow sense: the NIPS community sought to win credibility from the involvement of uniform perspectives in multiple fields rather than the incorporation of divergent views such as those held by Fodor, Pylyshyn, or McLaughlin. No matter, in 1988, Papert promised that the artificial neural network offered 'a vindication of behaviorism against Jean Piaget, Noam Chomsky, and all those students of mind who criticized the universalism inherent in behaviorism's *tabula rasa*'.[32] Notions of universality at NIPS thus persisted as the neural network remade behaviourist conceptions of mind, enabling researchers to sidestep—through appeals to interdisciplinarity—challenges to the connectionist orthodoxy.

Three: Single male speaker

Hubris with reference to mind-computer metaphors was not the only source of unity for participants at the Denver event. NIPS conference papers, the overwhelming majority of which were written by white, male authors working in academic and corporate computer science labs, represent the way in which neural networks remade a form of behaviourism that contained perspectives of affluence, whiteness, and masculinity. Writing in a conference paper, University of Cambridge researchers Anthony Robinson and Frank Fallside explored the application of neural networks to speech coding in which the authors trained a network on 40 seconds of speech from a single male speaker, using the results to outline three competing architectures for speech coding processes.[33] In research presented at the conference, Bell Labs scientist David Burr boasted of 'high accuracy exceeding 98%' on speech recognition problems. In Burr's research, all training data were provided, again, by a single male speaker.[34] Tony Materna, an executive at NIPS sponsor Hecht-Nielsen Neurocomputers, underscored the way in which such biases might collide with society when he told the *Schenectady Gazette* in 1987 'In business, in the near future, you'll simply have to speak into a word processing unit and it will write letters and reports based on what is spoken into the machine'.[35] Comments by Materna and research by Robinson, Fallside, and Burr draw attention to the way in which values are encoded into the network and, upon the release of products based on such systems, into society. In this way, NIPS provides a window into the way in which a failure to consider the identity of those conducting or participating in research—in this instance, white, male authors working in computer science labs—led to the incorporation of perspectives of masculinity, whiteness, and affluence within the bounds of neural networks.

---

[32] Seymour Papert, "One AI or Many?" 267.
[33] Anthony Robinson and Frank Fallside, "Static and Dynamic Error Propagation Networks with Application to Speech Coding," in "Proceedings," 639.
[34] David Burr, "Speech Recognition Experiments with Perceptrons," in "Proceedings," 144–49.
[35] Delthia Ricks, "Computer Researchers Look to Human Brain," 20.

Historians of computer history argue that the choice of the 'problem' that a system aims to solve, as well the raw materials on which such models are trained, governs which values and ideals are promoted at the expense of others. Nathan Ensmenger, for instance, has shown that the ability to play chess was widely considered by AI researchers to be a strong indicator of general intelligence during the second half of the twentieth century.[36] One reason for this, Adam argues, is that AI researchers have consistently extrapolated from a 'bounded problem solving situation to make an important claim about the nature of general problem solving.'[37] These dynamics recurred at NIPS during its foundational year in 1987. In a paper submitted to the conference, NIPS 'interdisciplinary advisory committee' member Terrence Sejnowski demonstrated a system capable of playing the boardgame backgammon. In describing his research as relating to 'significant progress in the development of intelligent systems,' Sejnowski's paper underscored how a highly narrow notion of intelligence, namely that of a backgammon player, could be positioned as representative of more general notions of human intelligence.[38] As I have shown, these assumptions escaped conference-floor discussions at the Sheraton Tech Center. Sejnowski spoke to the *New York Times* in the weeks following NIPS to discuss NETtalk, a neural network that pronounced English words.[39] The system had been trained on '1,000 of the most commonly occurring words' that appeared in a 1967 research corpus from linguists at Brown University.[40] According to an updated version of the training data set published by the UCI Machine Learning Repository, NETtalk was trained on words that included 'negro', 'black', 'wife', and 'race'.[41] Far from a view from nowhere, Sejnowski's work serves to once again highlight the way in which research related to NIPS encoded biases and behaviours to reinforce perspectives of affluence, whiteness, and masculinity through a failure to consider the identity of the 'thinking' subject.

Four: Utopia of mind

Although the constitutive politics of neuroscience—especially in the context of the cognitive revolution in the United States during the 1950s—has been documented by scholars, a commensurate study of the artificial neural

---

[36]. Nathan Ensmenger, "Is chess the drosophila of artificial intelligence? A social history of an algorithm." *Social Studies of Science* 42 (1) 9.
[37] Alison Adam, "Artificial Knowing" 37, 93.
[38] Gerald Tesauro and Terrence Sejnowski, "A 'Neural' Network that Learns to Play Backgammon," in "Proceedings" 794–803.
[39]George Johnson, "A Machine Learns to Speak: Mimicking the Neurological Networks," *New York Times*, November 29, 1987 18.
[40] Terrence J. Sejnowski and Charles R. Rosenberg, "Parallel Networks that Learn to Pronounce English Text," *Complex Systems* 1 (1987): 145-168.
[41]Anon., "Connectionist Bench (Nettalk Corpus) Data Set," UCI Machine Learning Repository, last modified 13 January 2021. https://archive.ics.uci.edu/ml/datasets/Connectionist+Bench+(Nettalk+Corpus).

network has so far been overlooked.[42] Jamie Cohen-Cole, for example, demonstrates the way in which proponents of open-mindedness in the 1950s United States utilised the cognitive revolution to argue for the superiority of American liberalism in the Cold War.[43] He does not, however, consider analogous trends in the histories of AI. Having previously observed that Skinner's behaviourism emerged alongside his idealisations of social order—and the way in which connectionists remade behaviourism and blurred the boundaries between cognition and computing in the late 1980s—I consider, in my final section, the ways in which NIPS represented moves to describe society using neural networks. In what follows, I reimagine the view from nowhere as a tool for interpreting the constitutive nature of shared conceptions of knowledge. I show that, much as the cognitive revolution has been linked with notions of liberal open-mindedness during the early Cold War period, the remaking of behaviourism within the connectionist programme contained its own strand of utopian politics.

With regards to behaviourism and connectionism, I argue that shared conceptions of knowledge concealed shared social logics. As Steven Shapin and Simon Shaffer explain: 'Wherever and whenever groups of people come to agree about what knowledge is, they have practically and provisionally solved the problem of how to array and order themselves. To have knowledge is to belong to some sort of ordered life; to have some sort of ordered life is to have shared knowledge.'[44] Agreement about the physical processes underpinning knowledge creation and intelligence at NIPS teased answers to problems of social arrangement and order, which was represented by the language used by neural network researchers during and around the event that echoed that of behaviourists earlier in the twentieth century. Both, I argue, suggested that a harmonious society could be designed through the control of behaviour via the careful management of external stimuli. Most recently, Shoshana Zuboff has characterised this behaviourism-as-control-technique—at least in relation to computation—as 'instrumentarian', defining it as the 'instrumentation and instrumentalization of [human] behaviour for the purposes of modification, prediction, monetization, and control.'[45] As I have mentioned, in his 1948 novel *Walden Two* Skinner outlined a vision for a society in which 'we use our knowledge about human behavior to create a social environment in which we shall live productive and creative lives.'[46] To raise a harmonious society, Skinner wrote in 1971, 'what we need is a

[42] Jamie Cohen-Cole, *The Open Mind: Cold War Politics and the Sciences of Human Nature*. (Chicago: University of Chicago Press, 2014) 2-10. Margaret Boden, "Mind as Machine."
[43] Jamie Cohen-Cole, "The Open Mind" 7.
[44] Steven Shapin and Simon Schaffer, *Leviathan and the air-pump: Hobbes, Boyle, and the experimental life*. (Princeton: Princeton University Press, 1985) 5.
[45] Shoshana Zuboff, *The Age of Surveillance Capitalism*, (London: Profile Books, 2019) 332.
[46] B. F. Skinner, *Walden Two* (New York: Macmillan, 1949) xvi.

technology of behavior….comparable in power and precision to physical and biological technology'.[47] In what follows, I examine descriptions of the mind and of the neural network using the language of social organisation, which, in 1987, I argue extended Skinner's vision of a technology of behaviour.

Indulging metaphors of social organisation, an eager media seized on connectionist conceptions of mind, machine, and society orbiting the conference. *Pittsburgh Post-Gazette* journalist Henry W. Pierce wrote in May 1987 that 'the mind, explained psychologist Marvin Minsky, is really a hierarchy of levels of knowledge. It's like a city government with many offices and agencies that work together in committees but have no mayor'.[48] In reporting the work of Rumelhart and NIPS contributor and Carnegie-Mellon University computer scientist James McClelland, Pierce reported that 'no single input from one unit to another is strong enough to generate behaviour by itself, just as no single voter in a presidential election can decide who will get the office.'[49] In the months leading up to NIPS, Minsky published *Society of Mind* in which he recast cognitive systems as a society of individual processes.[50] The historian Jonnie Penn emphasises Minsky's interest in the isomorphic relationship between social dynamics and intelligent action, writing that Minsky aimed to replicate the 'cultural cues' that enabled humans to perform intelligent behaviours within machine intelligence.[51] For Minsky, Rumelhart, and Pierce, society described the dynamics of the mind; so too, however, did the mind—and by extension the neural network—describe the dynamics of society. Yet such societies existed in abstraction only. Minsky's *Society of Mind*, with its rational, predictable, and malleable agents, posited a highly idealised view of a classless, harmonious, and egalitarian social order. In echoing Skinner's technology of behaviour, the metaphors of social organisation used by Minsky, Rumelhart, and Pierce indicate that shared conceptions of knowledge—in this instance, the nature of the mind—were tied to notions of interdisciplinarity and constituted shared social logics. These perspectives shaped the values and ideals contained within neural network research presented at NIPS, diffusing into connectionist systems behaviourist understandings of mind, machine, and society. Though Adam demonstrates that a failure to consider the thinking subject enabled the values and ideals of researchers to infiltrate AI systems, I argue that a study of the Denver conference signals that alternative approaches are needed for

[47] B. F. Skinner, *Beyond Freedom and Dignity*, (London: Hackett Publishing, 1971) 5.
[48] Henry Peirce, "Meetings of the mind." 28.
[49] Henry Peirce, "Meetings of the mind." 28.
[50] Marvin Minsky, *Society of Mind (*London: Heinemann, 1987).
[51] Jonnie Penn, "Inventing Intelligence: On the History of Complex Information Processing and Artificial Intelligence in the United States in the Mid-Twentieth Century," (PhD Thesis, University of Cambridge, 2020) 180.

investigating the formation of the *shared* or collective values, ideals, and perspectives on which AI systems are based.

Conclusion: NIPS or NeurIPS?

This paper has explored the role of interdisciplinarity in the remaking of behaviourism by the connectionist branch of artificial intelligence in the United Sates in the late 1980s. It has applied, and built upon, Alison Adam's view from nowhere argument to surface three ways in which NIPS represented a remaking of behaviourism within the connectionist programme through the use of rhetorical and technical appeals to the universality of research. First, notions of interdisciplinarity embodied by the conference legitimised behaviourist theories of artificial neural networks and brain function challenged by experts in other domains. Second, research presented at the conference reflected the way in which neural networks encoded behaviours that privileged perspectives of masculinity, whiteness, and affluence. Third, the use of the language of neural networks by those associated with NIPS to describe systems of social organisation indicated the presence of social logics that echoed the utopian character of Skinnerian behaviourism.

NIPS rebranded as 'NeurIPS' in 2018 following a backlash regarding the acronym's misogynist connotations.[52] As I have demonstrated, such bias was by no means a new phenomenon. Neural network research presented at NIPS in 1987 represented both conceptions of whiteness, affluence, and masculinity and a highly specific form of interdisciplinarity, offering an alternative perspective for observing the group that programme chairman Abu-Mostafa characterised as 'everyone'. I have shown that the view from nowhere allows for a consideration of the way in which the 'interdisciplinary interactions' of NIPS enabled mutually reinforcing conceptions of the brain and artificial neural networks, blending—with uncertain languages used in both press and publications—understandings of the bounds between behaviour and application, and between the mind and society. By deploying the view from nowhere to explore the relationship between shared conceptions of knowledge and shared social logics, I reimagined Adam's critique as a mechanism for appreciating a distinct strand of utopian politics within the connectionist turn. While Adam rightly identifies that a failure to consider the thinking subject enables the values and ideals of researchers to permeate their work, I argue that a new approach is needed for investigating the formation of the *shared* values on which AI systems are based, be they 'interdisciplinary' or otherwise. A

[52] Vishwam Sankaran, "'NIPS' AI conference renamed after 30 years over complaints of sexism," *The Next Web*, last modified 22 November 2018. https://thenextweb.com/artificial-intelligence/2018/11/22/nips-ai-conference-renamed-after-30-years-over-complaints-of-sexism/.

close study of NIPS in 1987 indicates that though a consideration of the thinking subject is a crucial first step in investigating the perspectives and ideals contained within artificial intelligence systems, scrutinising shared values demands a wider lens.

Bibliography

Primary Sources

—. “Abstract of Papers. IEEE Conf. on "Neural Information Processing Systems - Natural and Synthetic." Institute of Electrical and Electronics Engineers, June 30, 1989. https://apps.dtic.mil/dtic/tr/fulltext/u2/a212697.pdf.

—. “Advance Program.” *Institute of Electrical and Electronics Engineers,* November 1987. http://work.caltech.edu/nipsimages/sponsors.jpg.

—. “Connectionist Bench (NETtalk Corpus) Data Set.” *UCI Machine Learning Repository*. https://archive.ics.uci.edu/ml/datasets/Connectionist+Bench+(Nettalk+Corpus).

Abu-Mostafa, Yaser. “The first NIPS (now NeurIPS),” last modified 19 May 2020. http://work.caltech.edu/nips.

Ambros-Ingerson, J., Granger, R., Henry, H., and Lynch, G. “Partitioning of Sensory Data by a Cortical Network.” In *Proceedings of the 1987 International Conference on Neural Information Processing Systems*, 317–337 edited by Dana Anderson, New York: American Institute of Physics, 1988.

Burr, David. “Speech Recognition Experiments with Perceptrons.” In *Proceedings of the 1987 International Conference on Neural Information Processing Systems*, 144–153 edited by Dana Anderson, New York: American Institute of Physics, 1988.

CBS Tomorrow. "The Thinking Machine," broadcast 1961, MIT Museum Collections. https://techtv.mit.edu/videos/10268-the-thinking-machine-1961---mit-centennial-film.

Ersü, Enis and Tolle, Henning. “Hierarchical Learning Control - An Approach with Neuron-Like Associative Memories.” In *Proceedings of the 1987 International Conference on Neural Information Processing Systems*, 249–261 edited by Dana Anderson, New York: American Institute of Physics, 1988.

Fodor, J. and McLaughlin, B. P. “Connectionism and the problem of systematicity: Why Smolensky’s solution doesn’t work.” *Cognition*, 35(2):183–204. 1990.

Fodor, J. A. and Pylyshyn, Z. W. “Connectionism and cognitive architecture: A critical analysis.” *Cognition*, 28(1):3–71. (1988).

Hammerstrom, Dan. “The Connectivity Analysis of Simple Association.” In *Proceedings of the 1987 International Conference on Neural Information Processing Systems*, 338–347 edited by Dana Anderson, New York: American Institute of Physics, 1988.

Johnson, George. “A Machine Learns to Speak: Mimicking the Neurological Networks.” *New York Times,* November 29, 1987.

MacDonald, Bruce. “Connecting to the Past.” In *Proceedings of the 1987 International Conference on Neural Information Processing Systems*, 505–514 edited by Dana Anderson, New York: American Institute of Physics, 1988.

Mackie, Stuart., Graf, Hans., Schwartz, Daniel., and Denker, John. “Microelectronic Implementations of Connectionist Neural Networks.” In *Proceedings of the 1987 International Conference on Neural Information Processing Systems*, 515–523 edited by Dana Anderson, New York: American Institute of Physics, 1988.

Peirce, Henry. “Meetings of the mind.” *Pittsburgh Post-Gazette*, May 11, 1987.

Pollack, Andrew. “More Human Than Ever, Computer Is Learning to Learn.” *New York Times,* September 15, 1987.

Ricks, Delthia. “Computer Researchers Look to Human Brain.” *Schenectady Gazette,* March 5, 1987.

Robinson, Anthony and Fallside, Frank. “Static and Dynamic Error Propagation Networks with Application to Speech Coding.” In *Proceedings of the 1987 International Conference on Neural Information Processing Systems*, 632–641 edited by Dana Anderson, New York: American Institute of Physics, 1988.

Strother, Robert. “Thinking machines are getting smarter.” *Mechanix Illustrated*, (10):159–167, 1958.

Smolensky, Paul. “Analysis of Distributed Representation of Constituent Structure in Connectionist System.” In *Proceedings of the 1987 International Conference on Neural Information Processing Systems*, 730–739 edited by Dana Anderson, New York: American Institute of Physics, 1988.

Smolensky, P. “The constituent structure of connectionist mental states: A reply to Fodor and Pylyshyn”. *The Southern Journal of Philosophy*, 26(S1):137–161, 1988.

Sejnowski, Terrence J. and. Rosenberg, Charles R. “Parallel Networks that Learn to Pronounce English Text”, *Complex Systems* 1 (1987) 145-168.

Tesauro, Gerald and Sejnowski, Terrence J. “A 'Neural' Network that Learns to Play Backgammon.” In *Proceedings of the 1987 International Conference on Neural Information Processing Systems*, 794–803 edited by Dana Anderson, New York: American Institute of Physics, 1988.

Secondary Sources

Adam, Alison. *Artificial Knowing: Gender and the Thinking Machine*, Routledge: London, 1998.

Brown, Nik and Rappert, Brian. *Contested Futures: A Sociology of Prospective Techno-Science.* Oxford: Routledge, 2017.

Boden, Margaret. *Mind as Machine*. Oxford: Clarendon Press, 2006

Cave, Stephen and Dihal, Kanta. "The Whiteness of AI", *Philosophy and Technology*. 33, 685–703 (2020).

Cohen-Cole, Jamie. *The Open Mind: Cold War Politics and the Sciences of Human Nature*. Chicago: University of Chicago Press, 2014.

Dick, Stephanie. 'Of Models and Machines: Implementing Bounded Rationality'. *Isis* 106, no. 3 (September 2015): 623–34.

Ensmenger, Nathan. "Is chess the drosophila of artificial intelligence? A social history of an algorithm" *Social Studies of Science* 42 (1).

Ferrando, Francesca. "Is the post-human a post-woman? Cyborgs, robots, artificial intelligence and the futures of gender: a case study." *European Journal of Futures Research.* 43 (2014).

Jasanoff, Sheila and Sang-Hyun Kim. "Containing the Atom: Sociotechnical Imaginaries and Nuclear Power in the United States and South Korea." *Minerva* 47, no. 2 (June 2009): 119-146.

Jones, Matthew L. *Reckoning with Matter: Calculating Machines, Innovation, and Thinking about Thinking from Pascal to Babbage*. Chicago: The University of Chicago Press, 2016.

Minsky, Marvin. "Artificial Intelligence." *Scientific American*, September, 1966.

Minsky, Marvin. *Society of Mind.* London: Heinemann, 1987.

Nagel, Thomas. *The View from Nowhere*. Oxford: Oxford University Press, 1986.

Mikel Olazaran, "A Sociological Study of the Official History of the Perceptrons Controversy," *Social Studies of Science* 26, no. 3 (August 1996): 611-659.

Olazaran, M. *Historical Sociology of Neural Network Research*, PhD Thesis; University of Edinburgh, 1991.

Papert, S. "One AI or Many?" In Graubard, S., editor, *The Artificial Intelligence Debate: False Starts, Real Foundations*, 241–267 Cambridge: MIT Press, 1988.

Penn, Jonnie. "Inventing Intelligence: On the History of Complex Information Processing and Artificial Intelligence in the United States in the Mid-Twentieth Century." PhD Thesis, University of Cambridge, 2020.

Plasek, Aaron. "On the Cruelty of Really Writing a History of Machine Learning." in *IEEE Annals of the History of Computing*, vol. 38, no. 4 Oct.-Dec. 2016.

Rappa, M. A., & Debackere, K. *The emergence of a new technology: the case of neural networks.* Massachusetts Institute of Technology, Alfred P. Sloan School of Management, 1989.

Rappa, M. A., & Debackere, K *"*International survey on the neural network research community," Massachusetts Institute of Technology, Alfred P. Sloan School of Management, 1990.

Sankaran, Vishwam. "NIPS' AI conference renamed after 30 years over complaints of sexism", *The Next Web,* last modified 22 November 2018. https://thenextweb.com/artificial-intelligence/2018/11/22/nips-ai-conference-renamed-after-30-years-over-complaints-of-sexism/.

Suchman, Lucille Alice. *Human-Machine Reconfigurations: Plans and Situated Actions*. 2nd ed. Cambridge; New York: Cambridge University Press, 2007.

Schwartz, T. J. "1987, the Neural Year in Review." *Synapse Connection* (February), 1and 11-12, 1988.

Shapin, Steven., & Schaffer, Simon. *Leviathan and the air-pump: Hobbes, Boyle, and the experimental life.* New Jersey, Princeton University Press, 1985.

Skinner, B. F. *Beyond Freedom and Dignity.* London: Hackett Publishing, 1971.

Skinner, B. F. *Walden Two.* New York: Macmillan, 1976.

Zuboff, Shoshana. *The Age of Surveillance Capitalism*. London: Profile Books, 2019.