

UCLA

UCLA Electronic Theses and Dissertations

Title

The Nonmachinables: Asymptotic Labor and the Political Economy of Contemporary Information-Processing Systems

Permalink

<https://escholarship.org/uc/item/0r78w2tn>

Author

Justie, Brian

Publication Date

2023

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

The Nonmachinables:
Asymptotic Labor and the Political Economy of
Contemporary Information-Processing Systems

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy
in Information Studies

by

Brian Justie

2023

© Copyright by

Brian Justie

2023

ABSTRACT OF THE DISSERTATION

The Nonmachinables:
Asymptotic Labor and the Political Economy of
Contemporary Information-Processing Systems

by

Brian Justie

Doctor of Philosophy in Information Studies

University of California, Los Angeles, 2023

Professor Sarah T. Roberts, Chair

This study introduces the concept of “asymptotic labor” through three case studies examining the political economy of contemporary information-processing systems. It aims to contribute additional historical and theoretical perspective to the strong foundation of existing critical research that has revealed the “hidden” human workforce that props up the vast digital infrastructure of artificial intelligence. The first chapter chronicles the design and development of CAPTCHA and reCAPTCHA, suggesting a critical periodization of these cybersecurity systems based on their different methods of validating human and nonhuman users. The evolution of CAPTCHA is deeply intertwined with the rapid ascendance of machine learning as the dominant form of artificial intelligence in the mid-2000s, and presages the emergent methods of value capture that undergird these data-

intensive systems. The next chapter builds on this latter premise, utilizing semiotics as a method for dissecting the mechanisms of meaning-making and value production at the core of these complex information-processing systems, as well as the ways that they have become implicated in a broader set of political-economic conditions. Finally, an ethnographic account of a specialized class of workers at the United States Postal Service ties together the more theoretically-laden arguments of the preceding chapters by demonstrating the social, political, and material implications that the development and implementation of these information-processing systems have on the dwindling number of humans that remain embroiled in their continuing operations.

The dissertation of Brian Justie is approved.

Christopher M. Kelty

Safiya U. Noble

Miriam Posner

Sarah T. Roberts, Committee Chair

University of California, Los Angeles

2023

For Anna.

Table of Contents

List of Figures	vii
Acknowledgments	viii
Vita	ix
Introduction	1
1. “Little History of CAPTCHA”	18
Little [internet] history	20
“...Fleeting and secret images...”	29
“...The associative mechanism...”	39
“...Smaller and smaller...”	52
2. “Indexical AI”	56
The Digital Index: A World Below	60
Net Semiotics	64
Indexical AI	75
The Shape of a Neural Net	80
Indexing Images	88
Indexing Language	98
Indexical Labor and Algorithmic Capital	111
3. “The Nonmachinables”	114
The Bureau of Hards	116
RIPS 122K	119
Crash Program	124
Processing Progress	129
Long-Term Temporary	133
Very Hard Tasks	137
Conclusion	142
Bibliography	146

List of Figures

Figure 1.	<i>Contemporary reCAPTCHA using Generative Universal Perturbation.</i>	20
Figure 2.	<i>Workflow and interface design from Digital Equipment Corporation patent application (Lillibridge et al., 2001).</i>	35
Figure 3.	<i>A convolutional neural network identifying a white dog.</i>	65
Figure 4.	<i>Diagram of a neural net.</i>	71
Figure 5.	<i>Machine Learning outputs rules.</i>	75
Figure 6.	<i>Logical circuits in a “nervous net” (McCulloch and Pitts, 1943).</i>	81
Figure 7.	<i>Solving for XOR with one hidden layer.</i>	82
Figure 8.	<i>Logic tables for OR and XOR functions.</i>	83
Figure 9.	<i>XOR with hidden layers.</i>	84
Figure 10.	<i>A neural net solving for symmetry.</i>	87
Figure 11.	<i>Images produced using a generative adversarial neural network.</i>	89
Figure 12.	<i>Diagram of Fukushima’s neocognitron, consisting of alternating layers of S-cells and C-cells.</i>	92
Figure 13.	<i>Feature maps produced by each layer of LeNet.</i>	94
Figure 14.	<i>Architecture of the first implementation of LeNet, illustrating the relationship between convolutional and pooling layers.</i>	97
Figure 15.	<i>The Transformer architecture.</i>	104
Figure 16.	<i>Multihead attention.</i>	104
Figure 17.	<i>Softmax in the attention equation.</i>	105
Figure 18.	<i>Softmax in the attention mechanism.</i>	106
Figure 19.	<i>An article generated by GPT-2 in the style of science journalism.</i>	109
Figure 20.	<i>Remote Encoding Center, Salt Lake City, UT.</i>	122

Acknowledgements

This dissertation came together slowly and then all at once, and only because of the unwavering support, encouragement, and determination of Sarah T. Roberts. I cannot thank Sarah enough for illuminating a path forward, after a prolonged period of apathy and resignation. Likewise, I am deeply grateful for the generosity and insightfulness of my committee: Miriam Posner, Safiya Noble, and Chris Kelty.

Leif Weatherby helped set me on this fateful course almost seven years ago, and ultimately played a critical role in this project, as a vital collaborator in long-distance writing and, lately, long-distance running.

Yair Agmon has been a peerless friend and tireless peer throughout this entire process. Ulysses Pascal, Sam Kellogg, Kerry Keith, Kelly Steinmetz and many other comrades graciously read, edited, and critiqued components of this dissertation.

Tia Koonse and Saba Waheed, and the rest of my colleagues at the UCLA Labor Center, have collectively shifted my trajectory as a scholar in a profound way over the last three years, and it is truly an honor to be joining their ranks more formally in the days ahead.

My parents, Julie and Kevin Justie, have long provided me with the requisite space and support to be creative and curious throughout my life – a fact that I think I largely took for granted, and only retrospectively have begun to fully process and recognize.

Finally, I have the deepest admiration and gratitude for Anna McGovern, without whom this undertaking would not have been possible.

Chapter 1 is a version of “Little History of CAPTCHA,” published in *Internet Histories* in 2020. Two anonymous reviewers provided concise and invaluable feedback.

Chapter 2 is a version of “Indexical AI,” published in *Critical Inquiry* in 2022, in collaboration with Leif Weatherby, with whom the research, writing, and editing duties were equally shared.

Chapter 3 is a version of “The Nonmachinables,” published in *Logic* in 2021, benefiting greatly from the clear-eyed editing prowess of Alex Blasdel.

Vita

EDUCATION

- 2023 [anticipated] Ph.D.
Department of Information Studies
University of California, Los Angeles
Los Angeles, CA
- 2018 M.A.
Department of Media, Culture, and Communication
New York University
New York, NY
- 2012 B.S.
College of Media, Arts, and Design
Northeastern University
Boston, MA

PUBLICATIONS

- Justie, B.**, Koonse, T., Macias, M., Ray, J., Waheed S. (January 2022). “Fast-Food Frontline: COVID-19 and Working Conditions in Los Angeles.” *UCLA Labor Center*.
<https://www.labor.ucla.edu/publication/fast-food-frontline-covid-19-and-working-conditions-in-los-angeles/>
- Weatherby, L. A., **Justie, B.** (January 2022). “Indexical AI.” *Critical Inquiry*, 48(2), pp. 381-415.
<https://doi.org/10.1086/717312>
- Justie, B.** (17 May, 2021). “The Nonmachinables.” *Logic Magazine*.
- Justie, B.** (November 2020). Review of *Documentarity: Evidence, Ontology, and Inscription*, by Ronald E. Day. *Information & Culture* 55(3). https://infoculturejournal.org/book_reviews/day_justie
- Justie, B.** (November 2020). “Little History of CAPTCHA.” *Internet Histories* (5)1, pp. 30-47.
<https://doi.org/10.1080/24701475.2020.1831197>
- Herrera, L., **Justie, B.**, Koonse, T., Waheed, S. (October 2020). “Worker Ownership, COVID-19, and the Future of the Gig Economy.” *UCLA Labor Center*.
<https://www.labor.ucla.edu/publication/gigworkers/>

Justie, B. (21 September, 2020). “Bot or Not.” *Real Life*.

Justie, B. (11 March, 2020). “The Rise of AI-Capitalism: An Interview with Nick Dyer-Witheford.” *Los Angeles Review of Books*.

Justie, B. (August 2019). Review of *Inhuman Power: Artificial Intelligence and the Future of Capitalism*, by Dyer-Witheford et al. *Information, Communication & Society* 23(1), pp. 151-154.
<https://doi.org/10.1080/1369118X.2019.1651372>

Justie, B. (13 May, 2019). “Fuzzy Logic, Fuzzy Ethics.” *Real Life*.

Justie, B. (12 August, 2017). “Knock the Hustle.” *Jacobin*.

PRESENTATIONS

Weatherby L. A., **Justie, B.** (17 June, 2021). “Indexical AI.” *The Philosophy and Critical Thinking of AI*. The Oxford Research Centre in the Humanities, University of Oxford, UK.

Justie, B. (14 April, 2021). “CAPTCHA, reCAPTCHA, and the Aesthetics of Platform Capitalism.” *Politics and Ethics of Platform Labour: Learning from Lived Experiences*. Centre for Research in the Arts, Social Sciences, and Humanities, University of Cambridge, UK.

Justie, B. (6 September, 2019). “Critique in the Age of ImageNet and AlexNet.” *The Society for Social Studies of Science (4S) Annual Meeting*. New Orleans, LA.

Justie, B. (13 April, 2019). “Fuzzy Logic, Fuzzy Ethics.” *Theorizing the Web*. The Museum of the Moving Image, Queens, NY.

Justie, B. (19 June, 2017). “Algorithms, Assemblages, Quantification, and Queer Theory.” *Connected Life Conference, Digital Inequality*. Oxford Internet Institute, Oxford, UK.

Justie, B. (19 April, 2017). “Dada Data.” *Theorizing the Web*. The Museum of the Moving Image, Queens, NY.

i just got word that theres a big problem at the factory and its called “The Boss”

—@dril_gpt2¹

Introduction

“Years and years and years...”

“...And it’s getting worse and worse!”

These were the sobering retorts I received from two workers after asking about the state of technological progress in their ultra-specialized line of employment as “keyers” for the United States Postal Service. Each originally hired as entry-level staffers more than three decades ago, these two postal workers have defied the odds. On their very first day on the job, they were repeatedly made aware that the positions would be strictly temporary, soon to be made redundant thanks to rapidly advancing technological innovations in automated mail sorting. In the meantime, however, their job as keyers was to assist the existing automated systems tasked with reading the addresses on letters and packages. The machine readers of the late 1980s and early 1990s – utilizing Optical Character

¹ The absurdist Twitter account @dril, beloved by fellow users for its decade-long oeuvre of esoteric and often politically combative quips, was co-opted as training data by an anonymous user in 2019 and fed into GPT-2, a breakthrough machine-learning language model. This auto-generated tweet, effectively cribbing @dril's style, serves as a meta-critique of these systems, and resonates with Trebor Scholz's argument, cited below, about the “old” dynamics of political economy that seem to reliably reproduce themselves despite the emergence of “new” technologies. @dril_gpt2. Twitter Post. January 27, 2020, 3:03 a.m. https://twitter.com/dril_gpt2/status/1221750499688550406.

Recognition systems powered by early forms of machine learning – were impressive, but still rejected upwards of 10% of all addresses as illegible. Keyers were assigned the cleanup job, deciphering the remainder of errant letters and packages, ensuring they would still arrive at their intended destination.

By the late 1990s, there were more than 30,000 keyers stationed at 55 “Remote Encoding Centers” across the country, reviewing and rerouting more than 25 billion pieces of mail annually. Since then, the number of keyers, RECS, and illegible parcels has plummeted. Today, there are just over one thousand keyers left, all working out of shabby cubicles in an antiquated warehouse in Utah, responsible for about a billion parcels per year. This latter figure – representing less than 1% of total annual mail volume – also reflects the fact that the error-rate of automated mail sorting technology has likewise plummeted in the intervening decades.

Nevertheless, the mission-critical work performed by keyers carries on. “Those last few percentage points,” my informant told me, “take years and years and years.” The technical landscape had evolved considerably, and advances in automated mail sorting had indeed reduced the overall demand for keyers, but this keyer remained largely unphased, knowing through experience that there will always be exceptional cases that vex even the most advanced machinery, demanding careful review by a well-honed human eye. After all, he'd

been through boom-and-bust tech cycles before, and had seen Postal Service leadership repeatedly fall prey to the beguiling promises of technological solutionism.

His experience as a keyer had taught him firsthand that the complicated puzzle of postal operations was primarily political in nature – a constellation of incentives and investments, proffered by policymakers and stakeholders who seemed to rarely prioritize the original Postal Service mandate of timely, reliable, and universal mail delivery.² That is to say, the longstanding challenges associated with automated mail processing could not simply be chalked up as a series of technical problems in need of technical solutions. His colleague's account of the gradually degrading quality of the work demonstrates what happens in this lurch of mismatched technical expectations and technical execution, and the ways that workers end up shuttled in to plug the holes and cover the gaps. This interplay – *the less work there is, the worse it gets* – points toward something like a convergence point of a number of political-economic variables present within the contemporary moment of digital information technology, accelerated by the unchecked

² This mandate, stretching back to the Postal Service Act of 1972, has been one of persistent contention (see, e.g. Fuller) and remains an issue receiving scrutiny and reconsideration to this day. Fuller, Wayne. *The American Mail: Enlarger of the Common Life*. Chicago: University of Chicago Press, 1972. “Reevaluating the Universal Service Obligation.” Office of Inspector General, United States Postal Service, May 6, 2020. <https://www.uspsoig.gov/sites/default/files/reports/2023-01/RISC-WP-20-004.pdf>.

ascendance of venture capital and the platform capitalism it underwrites, and an increasingly austere labor market overseen by policymakers all too eager to buy into narratives of unidirectional technological progress.

Historicizing and conceptualizing this juncture is the task of this dissertation. Across three case studies, I will trace the contours of a key concept that persists in the contemporary political economy of information-processing technology, that I have deemed “asymptotic labor.” Above all, it builds on a strong foundation of existing critical research that has revealed the seemingly *hidden* workforce that props up the vast digital infrastructure of “artificial intelligence.” While many examples of this work have been documented, there has been relatively little systematic inquiry into the *tendential* nature of this labor. By this, I mean that there appear to be consistently observable dynamics underlying the value production associated with this specific form of labor, whereby the work performed by humans to produce, collect, prepare, maintain, modify, validate and/or repair the datasets and models powering these AI systems is subject to two divergent rules:

1. The sheer quantity (i.e. *breadth*) of work tasked to humans seems to strictly decrease over time, ostensibly because of ongoing and linear technological progress in information-processing systems.

2. The qualitative significance (i.e. *depth*) of – or, relatedly, the relative value produced by – this labor seems to grow exponentially over time, indicating that this labor is increasingly critical to the regular operation of these systems.

Taken together, I am arguing that there exists a tendency that resembles the geometric function known as an “asymptote,” where a line continuously approaches a value without ever reaching it. The basis of this argument, as I will show in the subsequent case studies, is that strictly speaking, some remainder of human labor will *necessarily* exist within the circuit of AI development, implementation, and maintenance, even if this human labor appears to superficially tend toward total disappearance, whether that be quite literal (elimination of work) or more figurative (displacement and/or degradation of work).³ Asymptotic labor can serve as a heuristic for understanding the argument that human labor will never be fully excised from technical systems by dint of mere technical progress alone. Even as the remaining work appears to diminish in scope, on the one hand, and on the other hand, this same labor nevertheless proves increasingly critical for operations,

³ For accounts of both literal and figurative disappearance of information workers, see Roberts, Sarah T. *Behind the Screen: Content Moderation in the Shadows of Social Media*. Yale University Press, 2019. Gray, Mary L., and Siddharth Suri. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Eamon Dolan/Houghton Mifflin Harcourt, 2019. Huet, Ellen. “The Humans Hiding Behind the Chatbots.” *Bloomberg*, April 18, 2016. Fussell, Sidney. “Behind Every Robot Is a Human.” *The Atlantic*, April 15, 2019. Solon, Olivia. “The Rise of ‘Pseudo-AI’: How Tech Firms Quietly Use Humans to Do Bots’ Work.” *The Guardian*, July 6, 2018, sec. Technology. Chen, Angela. “Inmates in Finland Are Training AI as Part of Prison Labor.” *The Verge* (blog), March 28, 2019.

and therefore manifests as evermore valuable to those that own or control these complex information-processing systems.⁴

Many of these overriding questions about labor and technology are not new, even if the particular technologies under examination skew contemporary. Scholars have long sought to understand the ways that machines contribute to both the enrichment and the exploitation of workers. Charles Babbage, often cited as a pioneer of what could be understood as a form of pre-digital computing, described at length the processes through which manual labor could be reimagined in order to render it susceptible to replacement by machines.⁵ Picked up by subsequent scholars of modern technology, this idea of “de-

⁴ Broadly, this assertion is informed by the claim that all technology is both deeply political and intrinsically social in its design and implementation, following influential scholars including Langdon Winner, Wiebe E. Bijker and Trevor Pinch, among many others. Pinch, Trevor J., and Wiebe E. Bijker. “The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other.” *Social Studies of Science* 14, no. 3 (1984): 399–441. Winner, Langdon. “Do Artifacts Have Politics?” *Daedalus* 109, no. 1 (1980): 121–36.

⁵ Marx’s famous account of capital’s deployment of machinery takes direct inspiration from Babbage. Marx differentiated the ‘machine’ from its predecessor – the mere ‘tool’ – and situated it as one of the primary catalysts of the rapid and tumultuous brought about by the Industrial Revolution. The capitalist’s newfound ability to capture and crystallize the ‘variable capital’ of labor power in the form of machinery as ‘fixed capital’ further deepened the divide between classes. Considerable scholarship has been dedicated to the question of whether or not Marx’s accounting of the human-machine relationship under capitalism has any bearing on the contemporary landscape. For a brief survey of these issues, see my review of Dyer-Witheford et al’s *Inhuman Power*. Justie, Brian. “Review: Inhuman Power: Artificial Intelligence and the Future of Capitalism.” *Information, Communication & Society* 23, no. 1 (January 2, 2020): 151–54. <https://doi.org/10.1080/1369118X.2019.1651372>.

skilling” holds considerable purchase for understanding the relationship between labor and technology, spanning a century of political-economic transformation, from Taylorist ‘scientific management’ of factory production to midcentury and contemporary forms of information work. Harry Braverman systematized this method of critique, demonstrating how manual labor came to be differentiated from so-called “knowledge” work, ultimately creating new means of extracting surplus value and, more broadly, expanding and deepening the grip of capital on an ever-growing working class.⁶

Feminist scholars of political economy and STS, including Lorraine Daston, Cynthia Cockburn, and Ruth Cowan, further problematized this line of thinking, showing how even the baseline distinction between manual labor and knowledge work failed to fully acknowledge the ways in which work and the (re)production of value are intrinsically gendered.⁷ It is this work, by and large, that formed the basis for much of the contemporary scholarship aimed at uncovering and foregrounding the many forms of otherwise overlooked, unheralded, and thankless work at the root of many complex technical systems. Indeed, the critical heuristic of discovering, documenting, and exposing

⁶ Braverman, Harry. *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. 25th Anniversary ed. edition. New York: Monthly Review Press, 1998.

⁷ Daston, Lorraine. “Calculation and the Division of Labor, 1750-1950.” *Bulletin of the German Historical Institute* 62, no. Spring (2018): 9–30. Cowan, Ruth Schwartz. *More Work For Mother*. Basic Books, 1983. Cockburn, Cynthia. *Machinery of Dominance: Women, Men and Technical Know-How*. London; Dover, N.H: Pluto Press, 1987.

“hidden labor” has been paramount in scholarship addressing the political economy of information-processing technologies.⁸ Feminist scholars of science and technology have led the way in this regard.⁹ As Jennifer S. Light has forcefully argued, the history of computing is a story whose arc is defined by that which it has chosen to diligently erase and omit.¹⁰ The countless women employed during wartime, providing labor as technicians and programmers that were demonstrably vital to war efforts were, nevertheless, repeatedly “rendered invisible.”¹¹

In some instances, this workforce was intentionally congealed into an homogenous mass, erasing the unique and varied contributions of individual women, a discursive practice in

⁸ Blok, Aad, Greg Downey, and Senior Lecturer Greg Downey. *Uncovering Labour in Information Revolutions, 1750-2000: Volume 11*. Cambridge University Press, 2003.

⁹ Abbate, Janet. “The Pleasure Paradox,” 211–27, 2010. <https://doi.org/10.1002/9780470619926.ch10>. Cowan, Ruth Schwartz. *More Work For Mother: The Ironies Of Household Technology From The Open Hearth To The Microwave*. Basic Books, 1985. Daston, Lorraine. “Calculation and the Division of Labor, 1750-1950.” *Bulletin of the German Historical Institute* 62, no. Spring (2018): 9–30. Green, Venus. *Race on the Line: Gender, Labor, and Technology in the Bell System, 1880-1980*. Duke University Press, 2001. Hicks, Marie. *Programmed Inequality: How Britain Discarded Women Technologists and Lost Its Edge in Computing*. MIT Press, 2017. Irani, Lilly. *Chasing Innovation: Making Entrepreneurial Citizens in Modern India*. Princeton University Press, 2019. Light, Jennifer S. “When Computers Were Women.” *Technology and Culture* 40, no. 3 (1999): 455–83. Roberts, Sarah T. “Commercial Content Moderation: Digital Laborers’ Dirty Work.” *Dirty Work*, 2016, 12. Suchman, Lucy. “Making Work Visible.” *Communications of the ACM* 38, no. 9 (September 1, 1995): 56–64. <https://doi.org/10.1145/223248.223263>.

¹⁰ Light, “When Computers Were Women.”

¹¹ *Ibid.*, 455.

stark contrast to the recurring theme throughout the history of computing that elevates the supposedly singular contributions of so-called men of genius. In addition to this tendency, women technicians were often cast as mere synecdoche, diminutively described in terms of the machines they worked with, further displacing their individual agencies, as evidenced by the prominent cohorts of “scanner girls” and “ENIAC girls.”¹² Both of the preceding practices had the effect of making women workers appear interchangeable, their contributions systematically devalued, and their jobs as technicians diminished for supposedly requiring no skillfulness or ingenuity. Together, these factors coalesce to produce and reify what has been identified as “feminized” labor, a phenomenon with roots much deeper than the burgeoning early days of digital computing.¹³ The process these scholars describe as “feminization,” it will be demonstrated, is central to producing and reifying the more recent phenomenon of asymptotic labor.

To this end, the feminized workers in Light’s historical account rarely were afforded public recognition as a highly specialized and vital element in the leading technology initiatives of the day. This fact is seemingly mirrored in the contemporary practices of asymptotic labor under investigation here – especially the work performed by data annotators in

¹² Ibid., 459.

¹³ See aforementioned Cowan (1985), Daston (2018), as well as Keilty, Patrick. “Tedious: Feminized Labor in Machine-Readable Cataloging.” *Feminist Media Studies* 18, no. 2 (March 4, 2018): 191–204. <https://doi.org/10.1080/14680777.2017.1308410>.

preparing datasets for the “indexical AI” discussed in Chapter 2 – wherein those performing the critical labor almost never appear in the vast technical literature published by AI researchers, and when they do it is only under the pretense that they are nameless cogs providing unskilled, rote, and clerical work, rather than agential subjects. The tasks often performed in roles subjected to the dynamics associated with asymptotic labor, in both their professional and more casual contexts, are understood to be little more than a temporary inefficiency – like the women programmers, called upon then subsequently cast aside, in Mar Hicks’ rewriting of the history of computing in Britain—that will eventually be overcome as contemporary AI systems gradually improve and become increasingly autonomous.¹⁴

While many journalists have by now taken stock of this phenomenon, the relative attention paid to the persistence of this “hidden” workforce performing mission-critical micro-tasks in the service of AI systems pales in comparison to the many breathless accounts by the popular press of innovation and progress in AI commercialization.¹⁵ As such, Light’s pioneering work provides a foundational reference for this project, which is similarly interested in exploring the ways in which work and workers are imagined and represented, if at all, by the academic computer scientists and industrial technologists

¹⁴ Hicks, *Programmed Inequality*.

¹⁵ Chen, “Inmates in Finland Are Training AI as Part of Prison Labor”; Fussell, “Behind Every Robot Is a Human”; Huet, “The Humans Hiding Behind the Chatbots”; Solon, “The Rise of ‘Pseudo-AI.’”

designing these systems, as will be unpacked in the deep reviews of technical literature in the subsequent case studies.

Accordingly, I am hopeful that the subsequent case studies presented here can serve as a minor contribution to the bountiful scholarship that is situated at this juncture of labor and technology, that makes a point of taking seriously the intricacies of both. Much of this scholarship has been undertaken, unsurprisingly, by a new school of feminist science, technology, media, and communications scholars, from which I draw much inspiration. This work is typified by, among others, Lilly Irani’s account of Mechanical Turk,¹⁶ Sarah T. Roberts’ deep ethnography of the traumatic labor of commercial content moderators,¹⁷ and Mary Gray’s revelatory work with Siddharth Suri on the sheer scale of hidden labor that undergirds contemporary technological infrastructure.¹⁸ While at times expanding well beyond the strict confines of information studies, Irani, Roberts, and Gray are exemplary waypoints in a long lineage documenting and assessing the multifaceted forms of labor done by information professionals, which appears to be especially susceptible to

¹⁶ Irani, Lilly. “Difference and Dependence among Digital Workers: The Case of Amazon Mechanical Turk.” *South Atlantic Quarterly* 114, no. 1 (2015): 225–34. <https://doi.org/10.1215/00382876-2831665>.

¹⁷ Roberts, Sarah T. *Behind the Screen: Content Moderation in the Shadows of Social Media*. Yale University Press, 2019.

¹⁸ Gray, Mary L., and Siddharth Suri. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Eamon Dolan/Houghton Mifflin Harcourt, 2019.

the *asymptotic* dynamics of simultaneous quantitative diminishment and qualitative degradation outlined here.

Over the past century, spanning Mary Salome Cutler Fairchild’s early accounts of women working in libraries,¹⁹ onward through to the numerous accounts of the particular precarities of librarianship in the 21st century,²⁰ this modality of labor and its many derivatives both in and outside of the confines of information studies, have trafficked under many names: knowledge work, information work, immaterial labor, affective labor, cognitive labor, digital labor, to name just a few.²¹ But as Trebor Scholz succinctly put it in the introduction to a sprawling compendium on digital labor: “These are new forms of labor but old forms of exploitation.”²² The phenomenon of asymptotic labor is, strictly

¹⁹ Fairchild, Mary Salome Cutler. “Women in American Libraries.” *Library Journal* 29 (1904): 157–62.

²⁰ Cope, Jonathan. “Neoliberalism and Library & Information Science: Using Karl Polanyi’s Fictitious Commodity as an Alternative to Neoliberal Conceptions of Information.” *Publications and Research*, January 1, 2014. https://academicworks.cuny.edu/si_pubs/4. Nicholson, Karen P. “The McDonaldization of Academic Libraries and the Values of Transformational Change.” *College & Research Libraries* 76, no. 3 (March 1, 2015): 328–38. <https://doi.org/10.5860/crl.76.3.328>.

²¹ For a thorough review of the similarities and differences of these variegated terms, see Wilkie (2011) and Dyer-Witheford et al (2019). Wilkie, Rob. “Global Networks and the Materiality of Immaterial Labor.” In *The Digital Condition*, 50–121. Class and Culture in the Information Network. Fordham University Press, 2011. <https://doi.org/10.2307/j.ctt14brzk3.6>. Dyer-Witheford, Nick, Atle Mikkola Kjosén, and James Steinhoff. *Inhuman Power: Artificial Intelligence and the Future of Capitalism*. London: Pluto Press, 2019.

²² Scholz, Trebor. *Digital Labor: The Internet as Playground and Factory*. Routledge, 2012. 1.

speaking, merely a symptom of the newfangled mechanisms of value production, capture, and capitalization – AI, writ large – made possible by these “old forms of exploitation.”

In brief, the three case studies comprising this dissertation will unfold as follows, each contributing toward the concretization of a working theory of asymptotic labor, with the intent to complicate and enrich the existing discourse surrounding AI and “hidden” labor:

1. An examination of CAPTCHA will serve to trouble the relationship between humans and machines in order to reveal a shifting value proposition associated with producing and maintaining this surprisingly slippery distinction. CAPTCHA and its many derivatives reveal the seemingly fleeting threshold between the interpretive abilities supposedly unique to humans and the ever-encroaching specter of digital mimicry. The distinction between humans and computers, this case studies demonstrates, cannot be understood strictly on ontological grounds, but rather is dependent upon the political-economic conditions under which both entities come into contact with one another. Ultimately, this study of CAPTCHA, reCAPTCHA, and their many derivatives effectively complicates the human/nonhuman binary, and raises questions about how emergent technology – namely, so-called “deep” machine learning – further imbricates this relationship and opens new doors for the forms of value production and capture that produce and accelerate the dynamics of asymptotic labor.

2. The next case study, which unpacks the technical innerworkings of “deep” machine learning systems, suggests one method by which to begin answering the preceding questions about human/nonhuman political economy. Drawing from critical semiotics, we outline how digital systems have become pliable and dynamic, appearing to perform feats of judgment, interpretation, and creativity that approximate human faculties in ways previously thought impossible. We attempt to grapple with the genuine novelties of these AI systems without also inadvertently reifying the bold claims of their boosters and benefactors who uncritically claim that the systems are actually producing forms of ‘intelligence’ identical to those possessed by humans. Instead, we demonstrate how the semiotic function of ‘indexicalization’ can produce something resembling human intelligence in a neural network through techniques of distributed representation whenever there is sufficient fidelity between the material human systems in which the network is embedded and the ultra-fine-grained representations of these systems that have been captured in the data used to train the networks. This feedback loop between human action and data capture, we argue, occupies a central position within the contemporary political economy of information-processing technology, and drives the tendency I have identified here as asymptotic labor. The “indexical AI” enabled by machine learning, unlike its predecessor “symbolic AI,” is extremely data-intensive, always requiring both *more* and *better* input data to render its human-

like faculties. The quest for *more* data entails further encroachments on personal privacy and autonomy, while the twin demand to procure *better* data requires additional oversight from humans.²³ Both incentives mean that humans will remain embroiled “in the loop” – and continuously subjected to the downward pressures of asymptotic labor – regardless of how fanciful and futuristic the narratives of progress AI appear to indicate.²⁴

3. Finally, having identified the moving parts of the equation of asymptotic labor in the first case study and modeled their interaction in the second, I dutifully enter the ‘hidden abode of production’ to observe asymptotic labor *in situ* at the United

²³ For additional historical perspective on the emergent conditions of surveillance eliciting “more” data, see Zuboff; and for a prescient account of how the actual data capture process can be designed to produce “better” data, see Agre (1994). Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs, 2019. Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. Public Affairs, 2019. Agre, Philip E. “Surveillance and Capture: Two Models of Privacy.” *The Information Society* 10, no. 2 (April 1994): 101–27. <https://doi.org/10.1080/01972243.1994.9960162>.

²⁴ “Humans in the loop” is a recurring idea within the development of computerized systems, and was the initial idea undergirding the conceptualization of early CAPTCHA predecessors (see, e.g. Baird & Popat, 2002). Additionally, the anthropology of technology Nick Seaver has repurposed the idea in a pithy way to argue that technical systems are never *merely* technical, writing that, “If you cannot see a human in the loop, then you just need to look for a bigger loop.” Baird, Henry S., and Kris Popat. “Human Interactive Proofs and Document Image Analysis.” In *Document Analysis Systems V*, edited by Daniel Lopresti, Jianying Hu, and Ramanujan Kashi, 2423:507–18. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002. https://doi.org/10.1007/3-540-45869-7_54. Seaver, Nick. “What Should an Anthropology of Algorithms Do?” *Cultural Anthropology* 33, no. 3 (August 21, 2018): 375–85. <https://doi.org/10.14506/ca33.3.04>.

States Postal Service.²⁵ Through an ethnographic study of a specialized set of USPS workers, I am able to ground the preceding argument about asymptotic labor in a concrete workplace setting that is itself the result of a tumultuous history of technological innovation and adaptation. Building upon the lessons drawn from my assessment of CAPTCHA and the ways that ‘humanness’ becomes implicated in the design of complex technical systems, these postal workers similarly are situated at the fleeting threshold of an always evolving automated system.²⁶ For more than three decades, these workers have been told their jobs as “keyers” were merely a temporary stopgap en route to full automation – making quite explicit the otherwise implicit tendencies associated with asymptotic labor. While the total number of keyers has shrunk considerably in the face of labor austerity and technical progress, those remaining in the position play the part of keystone in the foundation of postal operations. This case study reveals quite clearly, however, that

²⁵ After his prolonged critique of the analytical shortcomings of his contemporaries' political-economic thinking, Marx famously declared that one must follow the worker into the “hidden abode of production” in order to fully grasp the machinations of capital. Fraser, Nancy. “Behind Marx’s Hidden Abode.” *New Left Review*, no. 86 (April 1, 2014): 55–72.

²⁶ Philip Agre provides a critical methodology for making sense of how users – including workers – are imagined in the design of technical systems, which helped inform these case studies. Agre, Philip E. “Conceptions of the User in Computer Systems Design.” In *The Social and Interactional Dimensions of Human-Computer Interfaces*, 67–106. USA: Cambridge University Press, 1995.

asymptotic labor is not an inevitable fact of linear technical progress, but rather is the byproduct of convergent social, political, and material variables.²⁷

Across these case studies, three significantly different methodologies are utilized, offering multiple vantages from which to make sense of asymptotic labor and the political economy of contemporary information-processing technology.

²⁷ By “material variables” here, I refer primarily to the incredible growth in computing power realized over the past two decades. This phenomenon has been chronicled in depth by Tim Hwang. Hwang, Tim. “Computational Power and the Social Impact of Artificial Intelligence.” SSRN Scholarly Paper. Rochester, NY: Social Science Research Network, March 23, 2018. <https://papers.ssrn.com/abstract=3147971>.

1. “Little History of CAPTCHA”²⁸

In order to begin elaborating the concept of asymptotic labor, it is incumbent first to understand the evolving nature of the technology driving this tendency over the last two decades, as well as the political-economic conditions under which this technology has been developed and implemented. The mid-2000s re-emergence of machine learning, after a sustained period of dormancy, provides a relatively clear-cut line of demarcation in the development of CAPTCHA. Both CAPTCHA developers and the developers intent on cracking these puzzles began to experiment with machine learning during this period, presaging much of the subsequent decade in terms of both the computational techniques that would soon become dominant as well as the crowdsourced labor capture that have tended to accompany these technologies. In this chapter, I trace a periodization around this inflection point, naming the two eras of development “realist” and “relational,” pointing toward a significant difference in the methods of validation employed during each period of CAPTCHA implementation, each entailing distinctive arrangements of value production that portend the emerging trend of asymptotic labor.

²⁸ Justie, Brian. “Little History of CAPTCHA.” *Internet Histories* 5, no. 1 (January 2, 2021): 30–47. <https://doi.org/10.1080/24701475.2020.1831197>.

This article features primary historical research, drawing its data from a large cache of technical journals, standards documentation, patents, promotional materials and other 'gray' literature I gathered and compiled between 2017 and 2020. Additionally, I amassed a minor archive of visual ephemera associated with CAPTCHA, reCAPTCHA, and related cybersecurity initiatives. This includes a large number of images from blogs and message boards where users shared new or unusual instances of CAPTCHA they encountered while browsing the internet over the past two decades. Likewise, I made a habit of taking screenshots of all CAPTCHAs I encountered during day-to-day internet use, which made palpable the incredible frequency with which we are confronted by these puzzles, as well as the subtle changes they undergo that might otherwise be overlooked in the bluster of trying to solve the puzzle and advance as quickly as possible. Of note, in the time following publication of this article, for example, Google has begun to experiment with new visual puzzles that utilize newer machine learning techniques, including “generative adversarial perturbation,” which is a means of distorting the pixel values in an image such that a human sees a bit of uniform graininess, but another machine learning-based agent will be tricked into seeing something entirely different (Figure 1).²⁹ These techniques, still nascent, will be addressed more directly in Chapter 2.

²⁹ Zhang, Yang, Haichang Gao, Ge Pei, Shuai Kang, and Xin Zhou. “Effect of Adversarial Examples on the Robustness of CAPTCHA.” In *2018 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC)*, 1–109, 2018. <https://doi.org/10.1109/CyberC.2018.00013>.

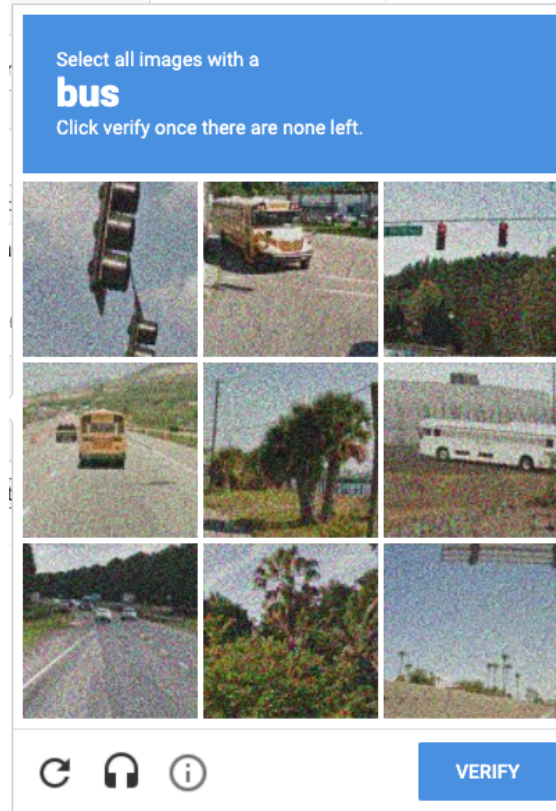


Figure 1. Contemporary reCAPTCHA using Generative Universal Perturbation.

Little [internet] history

“But mustn’t the photographer who is unable to read his own pictures be no less deemed an illiterate?”³⁰ This provocation appears at the end of an extended book review penned by Walter Benjamin in 1931, chronicling the near century-long arc spanning the invention

³⁰ Benjamin, Walter. *The Work of Art in the Age of Its Technological Reproducibility, and Other Writings on Media*. Harvard University Press, 2008.

of the daguerreotype up to and including nascent trends in contemporary art photography. What, Benjamin pondered, does a camera see that its operator cannot? Benjamin here poses a perennial question, one with intensifying stakes both then and now: are there inherent differences between the interpretative capacities of humans and of machines? For the last twenty-odd years, this question has been asked and answered hundreds of millions of times per day, each time the contemporary internet user is met with the infamous command to “Confirm Humanity.” As such, the interpretative exercises typified in CAPTCHA and reCAPTCHA have reduced the existentially-laden affirmation “*I’m not a robot*” to commonplace, phatic refrain. But to what end?

Benjamin’s meandering review—“Little History of Photography”—tracks shifting notions of mediation and materiality, laying out much of the conceptual groundwork that he returned to in the more widely known essay “The Work of Art in the Age of its Technological Reproducibility.”³¹ While the latter has been repurposed extensively as a theoretical guide for making sense of the digital and its discontents, the former has largely been overlooked in this respect. At stake in both essays, but perhaps more acutely in the “Little History,” were the lingering anxieties of modernity associated with automation and humanism. For Benjamin, this was precisely the nexus at which the new medium of photography productively intervened. Benjamin’s concluding question regarding illiteracy

³¹ Ibid, 19-55.

can be reread through this lens, as it immediately follows a reference to Eugene Atget's pioneering photographs of deserted Parisian streets. Benjamin sardonically advises that all of Atget's photographs be understood as the depiction of a crime scene—the incriminating act, of course, being the disappearance of the human, centering instead the apparatus and the environment rather than the photographer or the sitting subject.³²

This shift, for Benjamin, represented nothing less than a new paradigm wherein the photographic apparatus was taken to assert a newfound agency and autonomous existence, somehow indifferent to and in excess of human faculties. Hence Benjamin's seemingly accusatory stance regarding the contemporary photographer who has become unable to read, interpret, or identify the products of his own photographic practice. "The camera is getting *smaller and smaller*, ever readier to capture *fleeting and secret images*," Benjamin writes of the precipitous momentum of technological change, "whose shock effect paralyzes *the associative mechanisms* in the beholder."³³ Three things are evident in this claim: the first, an affirmation that the camera is capable of sui generis perceptive operations, catching glimpse of "fleeting and secret images" which elude the human interlocutor; second, that this capability causes a stultification in the viewer, whose capacity for interpretation is dependent on their ability to dynamically render meaningful associations

³² Ibid, 294.

³³ Ibid, 294, emphasis added.

between the captured representation and that which is represented; and third, that the photographic apparatus, like most technological innovations, becomes “smaller and smaller,” both in physical presence and in felt social imposition, as it becomes quotidian and backgrounded over time. The upshot, Benjamin seems to suggest, is that the ostensible monopoly on meaning-making heretofore held by humans might be fatefully jeopardized by advances in machinic mediation. To this end, Benjamin suggests that the textual “inscription”—or, in an earlier translation, the “caption”—may well become the most essential component of the photograph, inasmuch as it underwrites the human capacity for visual interpretation.³⁴

Our collective capacity for identifying and making sense of representational media—be it visual, aural, textual, or some combination thereof—and the cultural and political conditions in which these representations are understood to stand in relation to worldly referents, is an overriding concern for Benjamin. The “Little History,” therefore, might be reread as a critical historicization of representation as such, and the ways in which new modes of automated mediation tend to trouble the neatly articulated narratives of humanism long held dear. Nearly a century on from Benjamin’s initial writing, we find ourselves still confronted by an analogous angst. However the camera—Benjamin’s technical apparatus par excellence—has long since been eclipsed, largely subsumed by the

³⁴ Ibid, 195.

algorithmic automata and artificial intelligences lately occupying the role of anxiety-inducing apparatus *du jour*. Benjamin's handwringing over the emergent illiteracies signaled by his increasingly image-saturated era might indeed run parallel to the frequent and anguished invocations of chronic media illiteracy today.³⁵ Loudest among this discursive clamor is the omnipresent specter of malicious bots proffering spam, or "fake news," to online audiences unawares.³⁶ Once again, humanism and automation collide: what is at stake if computers can sufficiently mimic humans to such an extent that the two become indistinguishable?

It is my contention here that Benjamin's minor meditation on photography holds considerable purchase for those studying the contemporary internet, as well as the attendant concerns associated with automatic content identification, the topic to which this special issue is dedicated. This assertion is double-edged: automatic content identification comprises a suite of techniques that increasingly employ artificial intelligence, however a secondary task of perhaps equal or greater import, is the identification of content which has been produced by artificial intelligence, such that the latter might be appropriately regulated in the interest of human users. Put a bit

³⁵ boyd, danah. "Did Media Literacy Backfire?" *Points* (Data & Society) (blog), March 16, 2018. <https://points.datasociety.net/did-media-literacy-backfire-7418c084d88d>.

³⁶ Lazer, David, Matthew Baum, and Nicco Mele. "Combating Fake News: An Agenda for Research and Action." Harvard Kennedy School, May 2017. <https://shorensteincenter.org/combating-fake-news-agenda-for-research/>.

differently, and taking a page from Benjamin’s questioning of the photographer no longer able to interpret their own images, we might put forward an updated provocation regarding the contemporary machines able to generate content but unable to identify this selfsame content.

This latter conceit has served as the crux of a widely deployed method for combating spam and unwanted traffic online for more than two decades. CAPTCHAs, those pesky puzzles we are intermittently asked to solve while navigating the internet, are indeed little more than twofold content identification mechanisms: the user is asked to identify a seemingly trivial piece of content—most often words or images—in order for the computer to verify the identity of the user. Like Benjamin’s illiterate photographer, CAPTCHAs provide a constitutive case study of computers that are “unable to read [their] own pictures.” As this article will illustrate, such a premise opens onto a series of sociopolitical questions regarding the relationship between humans and machines—one of Benjamin’s principal concerns. The heuristic triad I have extracted from Benjamin’s assessment of the unsteady relation between humans and machines—viz. *“fleeting and secret images,”* *“the associative mechanism,”* *“smaller and smaller”*—shapes and reflects this article’s attempt to historicize the peculiar, yet illustrative curio of internet history that is CAPTCHA. If Benjamin’s “Little History” is to prove lucrative as a rubric for examining contemporary internet history, and moreover in excavating a genealogy of online content identification, it will arguably be along precisely these vectors.

As such, I will focus on the aforementioned and inherently redoubled nature of CAPTCHA as content identification system—or, to paraphrase Pierre Bourdieu, on the ways in which *identification identifies the identifier*.³⁷ What this means, more straightforwardly, is that by closely examining the subtle technical and social changes of these cybersecurity tools, we might likewise glean important insights regarding incipient modes of interaction between computers and humans, and the ways in which this relationship becomes codified in computational media and internet infrastructure. I will show how—and indeed why—CAPTCHA has gradually pivoted from what I am identifying as a *realist* framework of content identification toward a *relational* framework of content identification following its introduction in the late 1990s and steady rise to online ubiquity. This realist versus relational distinction calls to mind a long lineage of humanistic scholarship,³⁸ with

³⁷ Bourdieu, Pierre. *Distinction: A Social Critique of the Judgement of Taste*. Harvard University Press, 1984. 6.

³⁸ Realism and its ostensible opposite—whether termed relational, nominalist, relativist or otherwise—is, of course, a well-worn topic of critical inquiry. For additional perspectives that resonate with the conceptual issues signaled by CAPTCHA, see Hacking on representation and philosophy of science, Sekula on the political economy of photography, Winston on technology and aesthetics, Galloway on software and contemporary continental philosophy, or Burrell on digital infrastructure. The *realist* versus *relational* dichotomy is not cut-and-dry in the literature. Galloway and Burrell, for example, differ in their portrayal of actor-network theory, construing Bruno Latour as a realist or as a relational thinker, respectively. Hacking, Ian. *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge University Press, 1983. Sekula, Allan. “The Body and the Archive.” *October* 39 (1986): 3–64. <https://doi.org/10.2307/778312>. Winston, Brian. “A Mirror for Brunelleschi.” *Daedalus* 116, no. 3 (1987): 187–201. Galloway, Alexander. “The Poverty of Philosophy: Realism and Post-Fordism.” *Critical Inquiry*

particularly deep roots in science, technology, and infrastructure studies, but takes its foremost inspiration from Johanna Drucker’s critical reworking of the aesthetic foundations of data visualization practices.³⁹

It is worth noting at the outset that CAPTCHA—or, Completely Automated Public Turing Test for Telling Computers and Humans Apart—does not have a decisive origin story.⁴⁰ Rather, what began as a nonproprietary cryptography mechanism with contested

39, no. 2 (2013): 347–66. <https://doi.org/10.1086/668529>. Burrell, Jenna. “Thinking Relationally about Digital Inequality in Rural Regions of the U.S.” *First Monday* 23, no. 6 (June 1, 2018). <https://doi.org/10.5210/fm.v23i6.8376>.

³⁹ Drucker, Johanna. “Humanities Approaches to Graphical Display.” *Digital Humanities Quarterly* 005, no. 1 (March 10, 2011). Drucker, Johanna. *Graphesis: Visual Forms of Knowledge Production*. Harvard University Press, 2014.

⁴⁰ The evidence is inconclusive regarding the initial coining of the CAPTCHA acronym. Most popular references to CAPTCHA online date the acronym to 2003, while, in actuality, the term began to appear in the technical literature and conference proceedings several years prior. A 2001 paper by Baird et al., which appears to be the first published usage, also curiously contains a footnote citing both the captcha.net website and “personal communication” with Carnegie Mellon stakeholders regarding the CAPTCHA project, which date it back to 2000. The Internet Archive, however, has documentation of a live website at captcha.net beginning only in the fall of 2001, and ICANN’s domain registry confirms that this URL was first secured in February of that year, implying that the “personal communication” between Baird and CMU must have preceded the website launch, and therefore is the only, and ultimately unverifiable, source of this account of CAPTCHA’s apparent inception in 2000. Despite the contestation, what is clear, however, is that CAPTCHA’s rise to prominence, irrespective of the origin of its acronymic namesake is fully coincident with Luis von Ahn’s arrival at CMU in the fall of 2000, the same period in which Udi Manber, Chief Scientist at Yahoo!, had enlisted the computer science department to help reduce spam in his website’s chat rooms. Baird, Henry S., Allison L. Coates, and Richard J. Fateman. “Pessimial Print: A Reverse Turing Test.” In *Proceedings of the 6th International Conference on Document Analysis and Recognition*, 1154–

claims of inventorship, blossomed into a thriving community of diligent infosec researchers and resourceful hackers productively attempting to outwit one another and meaningfully advancing image recognition and document analysis technologies along the way, before ultimately ending up in the hands of a single technology conglomerate—Alphabet ne Google. For this reason, it is important to understand CAPTCHA as both a conceptual device which has remained relatively unchanged, insofar as it is still used for the identification of human users, but also as a variegated and constantly transforming collection of identification strategies pegged to evolving value propositions. A comprehensive genealogy of machines masquerading as humans (and vice versa) is far beyond the scope of this article. And despite the aforementioned acronym, such a story undoubtedly begins well before Alan Turing’s canonical test.⁴¹ For this reason, I will train my focus primarily on specific iterations and implementations of CAPTCHA between roughly 1996 and 2009, bridging important waypoints along its path of development in an effort to build toward an account of the implicit theories underlying and the sociopolitical stakes attending its use. To a large extent, this is how Benjamin approaches the topic of photography in the “Little History,” an admiral endeavor insofar as it leaves

58. Seattle, WA: IEEE Computer Society, 2001. Ahn, Luis von, Manuel Blum, and John Langford. “CMU-CS-02-117 (‘Telling Humans and Computers Apart (Automatically), or How Lazy Cryptographers Do AI’).” Technical Report. School of Computer Science, Carnegie Mellon University, February 2002. <http://reports-archive.adm.cs.cmu.edu/anon/2002/abstracts/02-117.html>.

⁴¹ Turing, A. M. “Computing Machinery and Intelligence.” *Mind*, no. 236 (October 1, 1950): 433–60. <https://doi.org/10.1093/mind/LIX.236.433>.

open unanswerable questions, while populating an historiographical map with generative clues and intriguing directives. I hope to achieve something similar in the pages that follow.

“ ... *Fleeting and secret images* ... “

Buried in the “Further Research” section, a mere two sentence coda tacked onto a preliminary draft, of an unpublished cryptography paper penned in 1996 by Moni Naor, is perhaps the first full explication of what eventually became known as CAPTCHA.⁴² Naor, an Israeli computer scientist with a penchant for visual cryptography, spent the preceding pages outlining a proposed fix for the growing issue of online spam gumming up free services like email and search.⁴³ The solution he outlined was centered around a clever repurposing of the Turing Test. Naor suggested that inquiring users be first presented with a content identification task—including “gender recognition” or “nudity detection” in images. These tasks, Naor demonstrated, might readily perplex the relatively

⁴² Naor, Moni. “Verification of a Human in the Loop or Identifications via the Turing Test.” Preliminary Draft. Weizmann Institute of Science, September 13, 1996.

⁴³ Brunton, Finn. *SPAM: A Shadow History of the Internet*. Infrastructures. MIT Press, 2013. Naor, Moni, and Adi Shamir. “Visual Cryptography.” In *Advances in Cryptology — EUROCRYPT’94*, edited by Alfredo De Santis, 1–12. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 1995. <https://doi.org/10.1007/BFb0053419>. Parikka, Jussi, and Tony D. Sampson. *The Spam Book: On Viruses, Porn, and Other Anomalies from the Dark Side of Digital Culture*. Hampton Press, 2009.

unsophisticated bots of the late 1990s, but would be “unambiguous” to human users.⁴⁴ Crucially, Naor suggested that these tasks be sourced from a large and predetermined database of questions and answers. Hence, the novelty of his final remarks, which notably step up the project’s ambit by reflecting on whether or not this entire process of “human-in-the-loop” verification might one day be fully automated. Naor’s hunch implies an eventual changing of guards: if sufficiently automated, the computer itself becomes the gatekeeper, no longer the mere administrator of a database already compiled and refined by humans. What’s more, Naor takes care to recommend that, if this automated identification process is to remain secure in the face of adversaries, it should make publicly available the program that is used to generate each test. The implication here—a profound one—is that security models that depend upon the withholding of key information are ultimately much less durable than models that prey on the ostensible differences in human and nonhuman interpretative capacity.

While the acronym “CAPTCHA” was not coined for another several years, Naor managed to touch upon nearly every one of its integral components in his draft’s concise coda: his ideal test would be *completely automated*, its code and data made *public*, conceptually modeled after the *Turing test*, and would serve as a reliable method for telling *computers and humans apart*. Naor did acknowledge, however, that there must remain “private

⁴⁴ Naor, 1996, 2.

random bits” of code inaccessible to the inquiring computer, lest the interrogator and interrogatee have parity access to the program, rendering the identification test altogether moot.⁴⁵ So while Naor may have dreamed up a completely automated process of verification, it could not, in fact, ever be completely public in his understanding. This is a vital characteristic of the fledgling infrastructure which would ultimately coalesce into what I’ve deemed to be CAPTCHA’s original realist framework, as it critically depends upon an initial concealment, or an essential withdrawal, as opposed to anything immanent to the content or the identification process per se. To solve one of Naor’s tests, according to this logic, would be nothing less than discovering a hidden, ontological truth. Computational randomness, therefore, serves as a functional surrogate for the *real*, as it presents a unidirectional mathematical process for computers, one vital to cryptography: very large prime integers generated randomly can be easily multiplied, but not easily factored. This approach merits its realist nomenclature because by withholding certain information from one subset of users—in this case, the “private random bits” used to match question and answer pairs in the readymade database—it implicitly reifies a sense of observer-independence implicit in the act of content identification. This suggests a real and unbreachable divide between humans and machines that a CAPTCHA-like test merely makes evident, always reproducing a simple binary outcome: bot or not. The content to be identified, then, appears as a series of “fleeting and secret images,” unknown

⁴⁵ Ibid, 4.

and unknowable to machines, but nevertheless real. In any case, however, there is no evidence that Naor actually attempted to build, let alone implement, a working version of this method himself.

Less than a month after Naor's draft was circulated, and with seemingly no knowledge of said draft, a team at Digital Equipment Corporation (D.E.C.) rolled out a similar deterrent to safeguard an online poll in advance of the impending presidential election.⁴⁶ By forcing users to first locate a slightly distorted image of the American flag that had been hidden somewhere on the polling webpage, its position randomly selected, the D.E.C. engineers hoped to stymie any attempts to vote more than once which would skew the poll results. This approach was easily circumvented with minor ingenuity and off-the-shelf programming, but did nevertheless have the intended effect of slowing down human attempts at ballot stuffing. In 1998, the following year, D.E.C. rolled out another, more sophisticated, tool for curbing undesirable interactions online. Alta Vista, D.E.C.'s flagship search engine product, was grappling with an increasingly unmanageable problem of infoglut. When a search query produced a large number of results, these results were returned in a ranked order based primarily on the number of pages linking to the queried term. But since Alta Vista maintained a full word index for all pages discoverable via its

⁴⁶ Lillibridge, Mark D., Martin Abadi, Krishna Bharat, and Andrei Z. Broder. *Method for selectively restricting access to computer systems*. United States US6195698B1, filed April 13, 1998, and issued February 27, 2001. <https://patents.google.com/patent/US6195698B1/en>.

search, clever users intent on boosting a certain search topic could easily exploit the participatory “ADD-URL” feature, which solicited input from users by allowing the manual indexing of new web pages. A simple automated script could submit innumerable new but functionally useless URLs and effectively co-opt search results.

A team of D.E.C. engineers, led by Alta Vista’s Chief Scientist Andrei Broder, was tasked with resolving this issue, and restoring order to search results. Finding a spark of inspiration in the user manual for their office scanner, which described in some detail its built-in Optical Character Recognition (OCR) feature,⁴⁷ Broder and his team developed a plan.⁴⁸ Even the most advanced OCR systems at the time, they realized, struggled to consistently “read” characters printed on physical documents that had been inadvertently stretched, rotated, or placed atop noisy background graphics.⁴⁹ For humans, these same distortions were trivial, hardly impinging at all upon overall legibility. This suggested to developers that there was a sizable gap between the human and the machine which could be capitalized upon. (The essentialism of perceptual faculties accorded to different types

⁴⁷ For two sharp accounts of the misgivings of OCR from the perspective of the critical humanities, see Cordell and Shoemaker. Cordell, Ryan. “‘Q i-Jtb the Raven’: Taking Dirty OCR Seriously.” *Book History* 20, no. 1 (2017): 188–225. <https://doi.org/10.1353/bh.2017.0006>. Shoemaker, Tyler. “Error Aligned.” *Textual Cultures: Texts, Contexts, Interpretation* 12, no. 1 (July 25, 2019): 155–82.

⁴⁸ Andrei Broder (Distinguished Scientist, Google), email message to Brian Justie, May 15, 2020.

⁴⁹ Baird, Henry S., and Kris Papat. “Human Interactive Proofs and Document Image Analysis.” In *Document Analysis Systems V*, edited by Daniel Lopresti, Jianying Hu, and Ramanujan Kashi, 2423:507–18. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002. https://doi.org/10.1007/3-540-45869-7_54. 509.

of users is yet another indication of the realist foundation underlying this approach.) While the randomness parameter in Naor's proposed test was intended to simply obfuscate the selection of predetermined question-answer pairs, Broder's solution was to randomize the test in full and automate its production: generate a random string of characters, twist and smush and elongate this string according to randomized variables, and randomly position this warped string within a randomized background pattern. This process produced an archetypal image, one now readily familiar to most internet denizens (Figure 2.) In order to access the "ADD-URL" feature, each Alta Vista user first had to spend several seconds deciphering one of these OCR-resistant puzzles. By 2001, "ADD-URL" spam had been reduced by 95%.⁵⁰

⁵⁰ Palo Alto Research Center. "Welcome to PARC's CAPTCHAs." CAPTCHA - History, February 28, 2003. <https://web.archive.org/web/20030407203015/http://www2.parc.com/istl/projects/captcha/history.htm>.

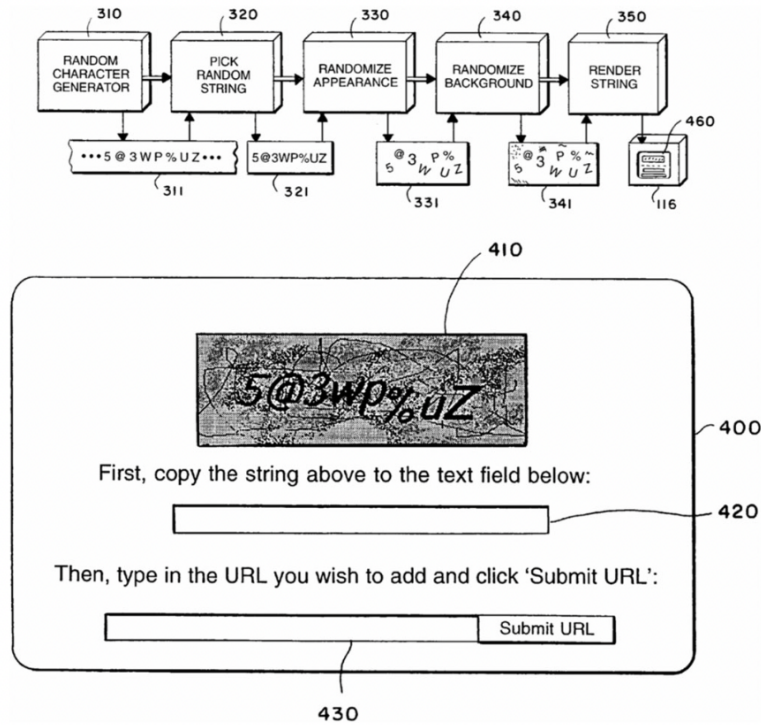


Figure 2. Workflow and interface design from Digital Equipment Corporation patent application (Lillibridge et al., 2001).

In the midst of this corporate anti-spam crusade, a small contingent of graduate students were waging a bot-war of their own. In 1999, slashdot, an online blog exceedingly popular with tech aficionados, ran an innocuous poll asking readers to cast their vote for the best computer science program in the country. Over several days, enterprising student engineers at the Massachusetts Institute of Technology (MIT) and Carnegie Mellon University (CMU) battled for the title, each unleashing automated voting scripts that evaded slashdot’s perfunctory IP-address checker. By the end of the polling period, MIT had narrowly edged out CMU by just 124 votes, each school receiving more than twenty

times the number of overall votes than the third-place school.⁵¹ Perhaps swayed by this promotional fanfare, a newly minted graduate of Duke University’s mathematics program enrolled in CMU’s computer science department that fall to pursue his doctorate. Over the next five years, under the advisement of the prolific cryptographer Manuel Blum (who had previously advised Naor’s dissertation), Luis von Ahn’s research would set into motion a gradual reshaping of the contemporary internet browsing experience. His manifold contributions to the burgeoning CAPTCHA project over this period helped to fundamentally tip the scales from the then dominant realist approach, toward the more dexterous, and ultimately pervasive, relational model.

Two papers coauthored by von Ahn in 2002 indicate the shifting terrain of CAPTCHA research, as well as the specific nature of his fledgling impact on the academic and commercial discourse. The first, published in September, and coauthored with two more senior scholars—von Ahn listed as third author, standard fare for an early graduate student—was a highly technical tract on the topic of steganography.⁵² A close cousin of classical cryptography, steganography differs from the former by focusing on the task of communicating private messages publicly, rather than on creating the ideal conditions for

⁵¹ Ahn, Luis von, Manuel Blum, Nicholas J. Hopper, and John Langford. “CAPTCHA: Using Hard AI Problems for Security.” In *Advances in Cryptology — EUROCRYPT 2003*, edited by Eli Biham, 2656:294–311. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003. https://doi.org/10.1007/3-540-39200-9_18.

⁵² See also Hopper, Nicholas, John Langford, and Luis von Ahn. “Provably Secure Steganography.” In *CRYPTO 2002: Advances in Cryptology*, 2002. https://doi.org/10.1007/3-540-45708-9_6.

outright concealment of a message’s content. In other words, the steganographer is interested in penning public messages with private meanings, wherein the message is available to all, but only a select few are aware that a secret has been transmitted. This premise arguably resonates with Benjamin’s investigation of the “fleeting and secret” messages that circulate below, or perhaps beyond, human faculties. Like the techniques proposed by Naor and Broder, the system outlined in this paper depended partially on the affordances of random number generation, stipulating that any “provably secure steganography” system was adequate only when its method of encryption rendered a message “computationally indistinguishable” from an arbitrary, or random, distribution. This is the computational basis upon which steganographic messages are made public to all, while operating privately for a specified subset. Put differently, the paper offered a formalized account of the ability to identify patterns where others see only randomness, or, in more humanist terms, to detect sense despite the appearance of nonsense. This paper does not address CAPTCHA explicitly, however it does ultimately run in parallel to the well-established orthodoxies in content identification-based online cybersecurity, which I’ve referred to here as realist.⁵³

⁵³ Baird, Henry S., Allison L. Coates, and Richard J. Fateman. “Pessimial Print: A Reverse Turing Test.” In *Proceedings of the 6th International Conference on Document Analysis and Recognition*, 1154–58. Seattle, WA: IEEE Computer Society, 2001. Chew, Monica, and Henry S. Baird. “BaffleText: A Human Interactive Proof.” In *Proceedings of Document Recognition and Retrieval X*, edited by Tapas Kanungo, Elisa H. Barney Smith, Jianying Hu, and Paul B. Kantor, 5010:305–16. Santa Clara, CA, 2003. <https://doi.org/10.1117/12.479682>. Chew, Monica, and J. D. Tygar. “Image Recognition CAPTCHAs.” In

Several months before this article appeared in the prominent journal *Advances in Cryptology*, however, an early draft of a technical report was circulated within CMU’s computer science community, with von Ahn occupying the first author position, a notable achievement for a second-year doctoral student. This report was momentous, subsequently appearing in conference proceedings in 2003 and eventually the *Communications of the Association for Computing Machinery* in 2004.⁵⁴ The foundation of von Ahn’s intervention, as first outlined in the 2002 CMU report, was to measure the efficacy of CAPTCHA against progress in artificial intelligence, a postulate signaled by the cheeky title: “Telling Humans and Computers Apart (Automatically), or How Lazy Cryptographers Do AI.”⁵⁵ It’s difficult to overstate how innovative this reframing proved

Information Security, edited by Kan Zhang and Yuliang Zheng, 3225:268–79. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004. https://doi.org/10.1007/978-3-540-30144-8_23. Hopper, Nicholas J., and Manuel Blum. “Secure Human Identification Protocols.” In *Advances in Cryptology — ASIACRYPT 2001*, edited by Colin Boyd, 2248:52–66. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001. https://doi.org/10.1007/3-540-45682-1_4.

⁵⁴ Ahn, Luis von, Manuel Blum, Nicholas Hopper, and John Langford. “Captcha: Telling Humans and Computers Apart Automatically.” In *Proceedings of Eurocrypt*, 294–311. Warsaw, Poland: Springer, 2003. https://doi.org/10.1007/3-540-39200-9_18. Ahn, Luis von, Manuel Blum, and John Langford. “Telling Humans and Computers Apart Automatically.” *Communications of the ACM* 47, no. 2 (February 1, 2004): 56–60. <https://doi.org/10.1145/966389.966390>.

⁵⁵ Ahn, Luis von, Manuel Blum, and John Langford. “CMU-CS-02-117 (‘Telling Humans and Computers Apart (Automatically), or How Lazy Cryptographers Do AI’).” Technical Report. School of Computer Science, Carnegie Mellon University, February 2002. <http://reports-archive.adm.cs.cmu.edu/anon/2002/abstracts/02-117.html>.

to be. Indeed, it is both the wellspring out of which CAPTCHA’s relational paradigm emerged, and a bellwether of the “deep learning” revolution in artificial intelligence that would crest over the subsequent decade, itself a relational alternative to the realist tradition of “symbolic AI.”⁵⁶

“ ... *The associative mechanism* ... “

The initial “Lazy Cryptographers” report is a scattershot five pages, beginning with a taxonomy of existing CAPTCHA types developed by the CMU team and concluding with a series of prescient speculations. Each of the four CAPTCHAs reviewed in its opening section typify the realist model, whereby something is withheld, or obscured by random-number generation, in order to ferret out nonhuman inquirers. “Gimpy,” the most familiar of these early CAPTCHAs, displayed seven words from the dictionary, selected at random, and distorted their shape and background; “Bongo” was a visual pattern recognition test,

⁵⁶ Cardon et al offer a thoroughgoing and nonlinear history of AI that acutely tracks this shift from realist to relational models of intelligence: “The symbolic approach that constituted the initial reference framework for AI was identified with orthodox cognitivism, in terms of which thinking consists of calculating symbols that have both a material reality and a semantic representation value. By contrast, the connectionist paradigm considers thinking to be similar to a massive parallel calculation of elementary functions – functions that will be distributed across a neural network – the meaningful behaviour of which only appears on the collective level as an emerging effect of the interactions produced by these elementary operations.” Cardon, Dominique, Jean-Philippe Cointet, and Antoine Mazières. “Neurons Spike Back (La Revanche Des Neurones: L’invention Des Machines Inductives et La Controverse de l’intelligence Artificielle).” *Réseaux* n° 211, no. 5 (2018): 173. <https://doi.org/10.3917/res.211.0173>. 4.

presenting two frames containing randomly generated and arranged shapes; “Pix,” an image recognition test, closely resembled the randomized question-and-answer method proposed by Naor; and “Eco,” the only nonvisual example outlined, in which the user was asked to transcribe a sequence of spoken characters and numbers that had been distorted, in effect producing an audio transposition of “Gimpy.”⁵⁷ The remainder of the report covered well-trodden ground, save for two particularly striking sections. These differed in tone and in scope from the rest of the report, and included the following provocations:

“Any program that passes the tests generated by a CAPTCHA can be used to solve a hard unsolved Artificial Intelligence (AI) problem.”

“Can we have a kind of SETI@home project in which web users help to classify all books in the library of congress?”

⁵⁷ von Ahn, Blum, and Langford, “CMU-CS-02-117.” Regarding accessibility, Matt May has written extensively on the issues posed by CAPTCHA, publishing an initial report in 2003 which has been continuously updated by May and collaborators in subsequent years. May, Matt. “Inaccessibility of CAPTCHA: Alternatives to Visual Turing Tests on the Web.” *Accessible Platform Architectures*. WAI Protocols and Formats Working Group, 2019 2003. <https://www.w3.org/TR/turingtest/#the-accessibility-challenge>.

The first explicitly links CAPTCHA to the ongoing concerns associated with AI research and development, a yet unacknowledged conjuncture. It also functionally shifts CAPTCHA away from a clearcut binary of success and failure—*did the bots get through or not?*—asserting instead a certain positive value proposition associated with the inevitable cracking of individual CAPTCHA puzzles. This conceit is evident in the report’s subtitle, suggesting that “lazy cryptographers” are, in fact, significantly contributing to progress in AI whenever their cybersecurity infrastructure fails. Second, by invoking SETI@home, a large-scale distributed computing project launched at Berkeley in 1999, von Ahn brings to the fore another latent value proposition underlying content identification-based CAPTCHAs.⁵⁸ The homebound “Search for Extra-Terrestrial Intelligence” ingeniously took advantage of the unused capacity of an enormous, linked network of personal computers repurposed for the computationally-demanding number-crunching required to process the unfathomably large datasets used in astronomy research. What von Ahn seems to be alluding to is a similar capitalization of surplus processing power, however this time not strictly limited to the voluntary provision of CPU leftovers, but rather to what he—perhaps unsettlingly—refers to here, and again elsewhere, as “stealing cycles from humans.”⁵⁹

⁵⁸ Korpela, E., D. Werthimer, D. Anderson, J. Cobb, and M. Leboisky. “SETI@home-Massively Distributed Computing for SETI.” *Computing in Science Engineering* 3, no. 1 (January 2001): 78–83. <https://doi.org/10.1109/5992.895191>.

⁵⁹ von Ahn, 2002, 5.

Before von Ahn gestured at the noble prospect of a crowdsourcing project to enrich the contents of the Library of Congress, as cited above, he offered a more brusque example of how this might work. Online pornography sites, long trafficking in the production and circulation of email spam, had lately encountered a stumbling block with the advent and adoption of CAPTCHA. But what if, von Ahn speculated, porn sites devised a procedure for rerouting each CAPTCHA encountered by one of their bots back to a human user elsewhere on the internet? Humans intent on accessing porn, for example, might make for a sizable and motivated demographic of puzzle-solvers, contributing a few seconds of their time in exchange for the desired content. A small price to pay for the individual user—a few extra seconds, another click or two—but, at the scale of many thousands of site visitors, an immensely valuable cache for the spammers.⁶⁰

Together, the two genuinely new contributions offered in this paper, yoking CAPTCHA to AI in a zero-sum contest, and the conceptual redefinition of a solved CAPTCHA as a “stolen human cycle,” provide the building blocks for the relational model which emerged

⁶⁰ Two years later, von Ahn’s prediction had seemingly come true, with spammers exchanging free porn for solved CAPTCHAs, as Cory Doctorow noted on the Boing Boing blog. Solving and creating captchas with free porn. The BBC reported a similar workaround in 2007, wherein hackers had further gamified the CAPTCHA-for-porn transaction. Doctorow, Cory. “Solving and Creating Captchas with Free Porn.” *Boing Boing* (blog), May 8, 2004. <https://boingboing.net/2004/01/27/solving-and-creating.html>. BBC News. “PC Stripper Helps Spam to Spread,” October 30, 2007. <http://news.bbc.co.uk/2/hi/technology/7067962.stm>.

over the next several years and remains in wide use today. The key to both of these suppositions, and their correspondence with the emerging relational paradigm, lies in a question raised by von Ahn towards the end of this report: “in general, how can we use humans to expand our computational abilities?”⁶¹ Humans and computers, this question seems to imply, are no longer ontological foes to be sorted by type, but rather relational compatriots, coequal participants in some greater project. If the realist paradigm presupposes the essentially distinct identities of humans and of computers, to be upheld and reified in the CAPTCHA encounter, the relational paradigm forces us to ask about the ramifications of a twofold expansion: computers, no doubt, expand the abilities of humans, but humans, von Ahn deftly illustrates, also can be put to use expanding the abilities of computers. Has von Ahn inadvertently furnished a critical insight long bandied about in science and technology studies, or does his decentering of the human point toward a fraught sociopolitical precipice? This line of inquiry becomes all the more weighty with the consideration of three subsequent moments in von Ahn’s budding career: the ESP Game, “Human Computation,” and reCAPTCHA.

⁶¹ von Ahn, 2002, 5.

The ESP Game, a collaboration with fellow CMU student Lauren Dabbish, was launched online in 2003 with an ambitious goal: to label every single image on the internet.⁶² These image labels were eminently valuable for text-to-speech accessibility online, for improved performance in image search and content moderation applications, and for compiling large datasets used to “train” machine learning systems. Von Ahn and Dabbish realized that if they could successfully gamify the menial task of assigning semantic labels to images, they could sidestep an otherwise cost-prohibitive process and expedite considerably the realization of their lofty goal. The game worked by displaying an image to two randomly paired players, prompting both to begin typing words to describe what they saw. With a countdown clock ticking away in the upper left-hand corner, the game placed users in a competitive environment in which each passing second meant a reduced point bounty, implicitly encouraging both players to submit as many potentially relevant words as quickly as possible. Points were awarded based on the number of words both players had “agreed upon” in the allotted time, with individual players climbing their way up the hotly contested game-wide leaderboard. Von Ahn and Dabbish, meanwhile, found themselves awarded one more cost-free, human-labeled image for their ever-expanding collection with each round of gameplay. Just four months after going live, the ESP game had been played

⁶² Ahn, Luis von, and Laura Dabbish. “Labeling Images with a Computer Game.” In *Proceedings of the 2004 Conference on Human Factors in Computing Systems - CHI '04*, 319–26. Vienna, Austria: ACM Press, 2004. <https://doi.org/10.1145/985692.985733>.

by 13,630 people, who had affixed 1,271,451 labels to 293,760 unique images.⁶³ By 2005, the number of labeled images had exceeded ten million.

While this game did not have any immediate application in the cybersecurity domain, it was von Ahn's first explication of and experimentation with content identification tasks of the relational variety. The image content successfully identified by players of the ESP Game was not verified according to the "private random bits" approach of standard cryptographic protocols, but rather was deemed to be accurate inasmuch as it manifested an index of social consensus. Moreover, this is also seemingly the first attempt to systematically crowdsource the labor-intensive task of constructing and annotating large image datasets for machine learning. This feat not only presages future iterations of CAPTCHA, but also predates Amazon Mechanical Turk, the platform typically associated with pioneering this type of clickwork, which launched in 2005.⁶⁴ We might reread this

⁶³ Ahn, Luis von. "Human Computation." Dissertation, Carnegie Mellon University, 2005. 27.

⁶⁴ Reports of so-called "CAPTCHA farms" began to emerge around 2008, which assembled large numbers of low-wage workers to solve CAPTCHAs on behalf of spammer and hacker clients. Marti Motoyama, likewise, has done extensive and invaluable research on the political-economic implications of this particular CAPTCHA-centric strand of outsourced labor. And for more on Mechanic Turk, see Irani. Danchev, Dancho. "Inside India's CAPTCHA Solving Economy." *ZDNet* (blog), August 29, 2008. <https://www.zdnet.com/article/inside-indias-captcha-solving-economy/>. Stone, Brad. "Breaking Google Captchas for Some Extra Cash." *Bits Blog* (blog), March 13, 2008. <https://bits.blogs.nytimes.com/2008/03/13/breaking-google-captchas-for-3-a-day/>. Motoyama, Marti, Kirill Levchenko, Chris Kanich, Damon McCoy, Geoffrey M Voelker, and Stefan Savage. "Re: CAPTCHAs – Understanding CAPTCHA-Solving Services in an Economic Context." *Proceedings of the USENIX*

innovation, according to Benjamin’s lexicon, as a progression from the ultimately futile attempt to produce and safeguard “fleeting and secret images,” like those contained in “Gimpy” and “Pix,” toward more durable techniques predicated on the transposition of the “associative mechanism[s]” of humans and computers. The latter, to reiterate, provides the conceptual basis for the relational model, foregrounding the ways in which we make and sustain meaningful connections between humans, computers, and the content either are capable of producing and processing.

In late 2005, von Ahn defended his dissertation, entitled “Human Computation.”⁶⁵ His jury included Jitendra Malik, a Berkeley computer scientist who had quite conclusively cracked the word-based “Gimpy” puzzles two years prior, catalyzing the shift away from realist models.⁶⁶ Recall, von Ahn had played a key role in developing “Gimpy” before turning toward the suite of more sophisticated and relational solutions outlined in his dissertation. Along with the ESP Game, “Human Computation” chronicled three other

Security, 2010. Motoyama, Marti. “Understanding the Role of Outsourced Labor in Web Service Abuse.” Dissertation, University of California, San Diego, 2011. <https://escholarship.org/content/qt87s441sw/qt87s441sw>. Irani, Lilly. “Difference and Dependence among Digital Workers: The Case of Amazon Mechanical Turk.” *South Atlantic Quarterly* 114, no. 1 (2015): 225–34. <https://doi.org/10.1215/00382876-2831665>.

⁶⁵ von Ahn, 2005.

⁶⁶ Mori, G., and J. Malik. “Recognizing Objects in Adversarial Clutter: Breaking a Visual CAPTCHA.” In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2003. Proceedings., I-134-I-141. Madison, WI, USA: IEEE Comput. Soc, 2003. <https://doi.org/10.1109/CVPR.2003.1211347>.

similar games, each of which cleverly disguised a piecemeal content identification task behind an addicting, gamified facade.⁶⁷ But for von Ahn, what he had developed throughout the course of his doctoral work, was not just an assortment of online games, but rather a generalizable method for “utilizing human processing power.”⁶⁸ This sentiment, first articulated as “stealing cycles from humans” in the “Lazy Cryptographers” report, took on a new sheen in the dissertation version, and was rebranded to avoid the extractive connotations as “games with a purpose.”⁶⁹ This rhetorical move has the effect of reinforcing a newfound discursive framework, in which users and computers coappear as esteemed colleagues in the ongoing production and maintenance of their shared online environment, a far cry from the more rigid approach of discrete classification which characterized the initial stages of CAPTCHA development. This recontextualization, spearheaded by von Ahn, helped to create the technical and social auspices under which CAPTCHA transmogrified from an academically-driven cybersecurity concern, to the inordinately valuable apparatus under proprietary control that is ubiquitous today. The transition was swift.

⁶⁷ Ahn, Luis von, Mihir Kedia, and Manuel Blum. “Verbosity: A Game for Collecting Common-Sense Facts.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '06*, 75. Montreal, Canada: ACM Press, 2006. <https://doi.org/10.1145/1124772.1124784>. Ahn, Luis von, Ruoran Liu, and Manuel Blum. “Peekaboom: A Game for Locating Objects in Images.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '06*, 55. Montreal, Canada: ACM Press, 2006. <https://doi.org/10.1145/1124772.1124782>.

⁶⁸ von Ahn, *Human Computation*, 11.

⁶⁹ *Ibid.*, 3.

In July of 2006, the search engine giant Google hosted von Ahn for an installment of the company’s “TechTalk” series, where he presented his doctoral work for a small crowd of executives and engineers.⁷⁰ The following month, Google announced that it had licensed the software undergirding the ESP Game, and soon after launched the Google Image Labeler. Google’s version was a near replica of its predecessor, dedicated primarily to improving the reliability of Google Image Search results. The allure of coordinating with a stranger online to produce a “winning” consensus, a tactic pioneered by the ESP Game, made for a thrilling gaming experience. Amplified by Google’s reach and resources, more than 200,000 players had contributed more than 50 million image labels to the company’s coffers by 2008.⁷¹

With the pipeline between von Ahn, a newly appointed associate professor at CMU, and Google now well established, stock can be taken of the third, and perhaps most impactful,

⁷⁰ *Human Computation*. Google TechTalk. Mountain View, CA, 2006. <https://www.youtube.com/watch?v=tx082gDwGcM>.

⁷¹ Compare the scale and speed of this achievement to the concurrent ImageNet project, a pathbreaking visual dataset for machine learning comprising 15 million images annotated by nearly 50,000 Mechanical Turk workers between 2007 and 2010. Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. “ImageNet: A Large-Scale Hierarchical Image Database.” In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–55, 2009. <https://doi.org/10.1109/CVPR.2009.5206848>. Ahn, Luis von, and Laura Dabbish. “Designing Games with a Purpose.” *Communications of the ACM* 51, no. 8 (August 1, 2008): 57. <https://doi.org/10.1145/1378704.1378719>.

of von Ahn’s relational revolution as alluded to above: reCAPTCHA. The inaugural website went live in the summer of 2007 with a galvanizing mandate: “STOP SPAM. READ BOOKS.” The CMU team behind this effort, led by von Ahn, sought to put the lessons of his dissertation into practice at unprecedented scale, presenting reCAPTCHA as a means by which to capitalize upon the more than 60 million CAPTCHAs that were being solved each day.

Where CAPTCHA generated squiggly text strings on command, reCAPTCHA sourced its textual input from the massive storehouses of typeset, printed material that were being digitized by the likes of the Google Books Project and the nonprofit Internet Archive.⁷² Nearly 20% of all the content in these text sources was unrecognizable by OCR systems, due to material degradations like smudges, tears, and waterlogs, as well as intentional typographic and graphic design idiosyncrasies which proved vexing to machine readers. While these “found” words still underwent random transformations upon presentation in a reCAPTCHA puzzle—lending them the signature warped aesthetic—the computer’s method of evaluation was no longer dependent upon randomization alone. Rather, by presenting the user with two distinct words, one of which had already been successfully digitized, the user’s input was now to be evaluated according to a surreptitious control

⁷² Ahn, Luis von, Benjamin Maurer, Colin McMillen, David Abraham, and Manuel Blum. “ReCAPTCHA: Human-Based Character Recognition via Web Security Measures.” *Science* 321, no. 5895 (September 12, 2008): 1465–68. <https://doi.org/10.1126/science.1160379>.

variable. In each reCAPTCHA, the computer presented a known element and an unknown element side by side to the user, but, unlike in CAPTCHA, the unknown element was genuinely unknown to the computer administering and evaluating the test.

This is a meaningful departure from the original CAPTCHA framework which utilized randomness in order to effectively alter the “real” identity of a known element. What this means is that reCAPTCHA operates according to an altogether different mode of content identification, wherein the relationality between two elements (known and unknown) was mediated by the relationality between multiple user interpretations of these elements. In other words, these tests recast content identification tasks as an interplay of the relation between relations. reCAPTCHA did not prepossess a “secret” and subsequently judge the submitted solutions for fidelity to a withheld “real” identity. Rather, it sought to productively convene multiple unwitting internet users, bringing them into contingent relation in order to identify content vis-a-vis consensus. When three subsequent users submitted the same solution for the unknown word, for example, it was temporarily added to the pool of known control words. Through this process, each word piped in from the OCR detritus was assigned a probabilistic weight that was progressively refined with each additional submitted solution. Words were determined to be “solved” for digitization purposes when they surpassed a given weight threshold.⁷³ If six users shown the same

⁷³ von Ahn et al., *reCAPTCHA*, 1466.

unknown word were unable to reach a consensus—which occurred about 4% of the time—it was discarded as “unreadable.”

In 2009, Google acquired reCAPTCHA for an unspecified amount, quickly spinning it out into an API made available to third-party developers.⁷⁴ Within two years, more than 200 million words per day—primarily sourced from the New York Times archive and the ever-expanding Google Books collection—were being deciphered, transcribed, and digitized by humans around the world using reCAPTCHA’s relational content identification infrastructure.⁷⁵ Von Ahn’s 2002 moonshot of a general purpose SETI@home had come to fruition, albeit with private beneficiaries at the helm, in lieu of public institutions like the Library of Congress as initially proposed.

⁷⁴ Ahn, Luis von. “Teaching Computers to Read: Google Acquires ReCAPTCHA.” *Official Google Blog* (blog), September 16, 2009. <https://googleblog.blogspot.com/2009/09/teaching-computers-to-read-google.html>. Wyszomierski, Michael. “Protect Your Site from Spammers with ReCAPTCHA.” *Official Google Webmaster Central Blog* (blog), January 26, 2010. <https://webmasters.googleblog.com/2010/01/protect-your-site-from-spammers-with.html>.

⁷⁵ That same year, von Ahn left Google to found Duolingo, the now extraordinarily popular language-learning and translation app, another example of its creators’ knack for gamification and crowdsourcing. Google invested \$45 million in von Ahn’s new endeavor the following year. Gugliotta, Guy. “Captchas Have Us Deciphering Old Text Through Woozy Web Clues.” *The New York Times*, March 28, 2011, sec. Science. <https://www.nytimes.com/2011/03/29/science/29recaptcha.html>.

“ ... Smaller and smaller ... ”

To conclude, I'd like to return to Benjamin's observation regarding the progressive diminution of the photographic apparatus, an intuition seemingly confirmed by the omnipresence of personal and surveillant cameras today. CAPTCHAs have likewise begun to fade into the background of the internet in recent years, ever intent on manifesting as a “smaller and smaller” imposition on the user, all the while ensuring its overall jurisdiction continues to expand. Following its inception and widespread adoption, reCAPTCHA puzzles evolved to include not just splotchy words pulled from old books, but also things like house numbers and dense streetscapes sourced from Google Maps. But in 2013, a new iteration was quietly rolled out called No CAPTCHA reCAPTCHA, which arguably cemented the relational paradigm of cybersecurity precisely by ditching the overt content identification tasks altogether.⁷⁶ Google's No CAPTCHA reCAPTCHA was built atop a new suite of “risk analysis” tools, which preemptively gathered information about the user, including their IP address and cached cookies, while simultaneously capturing behavioral information including cursor movements and mouse clicks made by the user before, during, and after any actual engagement with a CAPTCHA puzzle.

⁷⁶ Shet, Vinay. “ReCAPTCHA Just Got Easier (but Only If You're Human).” Google Online Security Blog (blog), October 25, 2013. <https://security.googleblog.com/2013/10/recaptcha-just-got-easier-but-only-if.html>.

Altogether, this method aims to gather in advance sufficient evidence to determine the identity of any user, human or otherwise. This meant that, moving forward, most users—excepting, of course, the “risky” ones—would no longer be forced to solve the familiar word- and image-based content identification puzzles. Rather, they would be simply prompted to click a checkbox emblazoned with the promissory caption: “I’m not a robot.” The content identified and the content identifier, under this guise, become coterminous: one’s identity functionally reduced to the ongoing production of identifiable content. Identity, thus, becomes something like an epiphenomenon which emerges out a constellation of discrete metrics measuring the behavior and performance of any one user against that of all other users. Google’s product manager for reCAPTCHA put it succinctly: “today the distorted letters serve less as a test of humanity and more as a medium of engagement to elicit a broad range of cues that characterize humans and bots.”⁷⁷ In 2017, Google began to experiment with a fully backgrounded version called Invisible reCAPTCHA, eliminating the checkbox altogether, rendering the whole of one’s internet environ an arena of relationalized cybersecurity.⁷⁸

Critical scholars of digital technology including Phil Agre and Colin Koopman offer highly theoretical accounts heralding the realization of precisely this model of data capture. Agre,

⁷⁷ Ibid.

⁷⁸ Google Developers. “Invisible ReCAPTCHA.” Accessed January 30, 2020. <https://developers.google.com/recaptcha/docs/invisible>.

well before it had become materially viable, describes a prospective mode of digital surveillance no longer premised on the collection of additional, discrete data points, but rather on the progressive remodeling of behavior.⁷⁹ Koopman, twenty-five years later, closes this loop by rehistoricizing the flattened distinction between ourselves and our data.⁸⁰ And finally, Drucker, as alluded to above, writing from the perspective of the digital humanities, has made the overarching point here all the more decisively, by reframing data as “capta,” a more generative term which takes as its starting point the fundamental constructedness and constitutional relationality of all data practices.⁸¹

The trajectory from CAPTCHA to reCAPTCHA to No CAPTCHA reCAPTCHA to Invisible reCAPTCHA is the crystallization of this telos. Recall, Walter Benjamin’s “Little History of Photography” traces the contours of photography’s incipience, revealing a patchwork of boosterism and skepticism—one mirrored in the turbulent uptake of today’s advancing technologies. Photography proved to be a revolutionary medium for Benjamin precisely inasmuch as it unsettled deeply held beliefs about what it meant to be human. One need not ascribe any radical potentiality to CAPTCHA or any of its derivatives,

⁷⁹ Agre, Philip E. “Surveillance and Capture: Two Models of Privacy.” *The Information Society* 10, no. 2 (April 1994): 101–27. <https://doi.org/10.1080/01972243.1994.9960162>.

⁸⁰ Koopman, Colin. *How We Became Our Data: A Genealogy of the Informational Person*. Chicago, IL: University of Chicago Press, 2019.

⁸¹ Drucker, 2011.

however, in order to discover within it a veritable palimpsest of human-computer relations, documenting an ongoing ebb of identification and disidentification.

2. “Indexical AI”⁸²

With the groundwork now laid by the preceding genealogy of CAPTCHA, we can move forward with a more analytical approach to understanding what – if anything – is truly novel about machine learning, and the means by which it has produced a set of political-economic conditions that give rise to asymptotic labor. In order to meaningfully contrast machine learning with its predecessor, “symbolic AI,” we employed a methodology of semiotic analysis in order to better understand the role of the “symbol” in the latter, and identify the key signifying function at play in the former, which we argue is the “index” – ergo, our neologism “indexical AI.”

Critical semiotics enjoyed its heyday as an analytical method more than half a century ago, at the time providing both terminology and tools for a class of structuralist theorists to systematize their claims about meaning-making, culture, and power.⁸³ However in the wake of poststructuralism, declarative accounts of sign and signifier gave way to less rigid and more relational theories of meaning, which undermined much of the analytical power

⁸² Weatherby, Leif, and Brian Justie. “Indexical AI.” *Critical Inquiry* 48, no. 2 (January 2022): 381–415. <https://doi.org/10.1086/717312>.

⁸³ Hawkes, Terence. *Structuralism & Semiotics*. University of California Press, 1977.

previously ascribed to semiotics.⁸⁴ Since then, semiotics has remained somewhat cloistered as a niche interest in certain corners of the academy, including especially the field of linguistic anthropology, which greatly influenced our examination of artificial intelligence.⁸⁵ Two kindred scholars in this domain, Paul Kockelman and Michael Castelle, have recently contributed semiotically-driven analyses to the growing discourse of critical AI studies, which resonate our argument about the emergence of “indexical AI.”⁸⁶

Elsewhere, a number of scholars have foregrounded the ways that semiotics can provide a new entry point for political-economic research, especially within regards to the question of how digital technology has changed traditional ideas about work and value.⁸⁷ This latter line of inquiry is of particular relevance here, as it provides a bridge between our

⁸⁴ Derrida, Jacques. “Structure, Sign, and Play in the Discourse of the Human Sciences.” In *The Structuralist Controversy: The Languages of Criticism and the Sciences of Man*, edited by Richard A. Macksey, 40th anniversary edition., 186–94. Baltimore, Md: Johns Hopkins University Press, 2007.

⁸⁵ Keane, Webb. “Semiotics and the Social Analysis of Material Things.” *Language & Communication* 23, no. 3–4 (July 2003): 409–25. [https://doi.org/10.1016/S0271-5309\(03\)00010-7](https://doi.org/10.1016/S0271-5309(03)00010-7). Gal, Susan, and Judith T. Irvine. *Signs of Difference: Language and Ideology in Social Life*. Cambridge: Cambridge University Press, 2019. <https://doi.org/10.1017/9781108649209>.

⁸⁶ Castelle, Michael. “The Social Lives of Generative Adversarial Networks.” In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 413. FAT* ’20. New York, NY, USA: Association for Computing Machinery, 2020. <https://doi.org/10.1145/3351095.3373156>.

Kockelman, Paul. “The Epistemic and Performative Dynamics of Machine Learning Praxis.” *Signs and Society* 8, no. 2 (March 2020): 319–55. <https://doi.org/10.1086/708249>.

⁸⁷ Gandini, Alessandro. “Digital Labour: An Empty Signifier?” *Media, Culture & Society*, August 12, 2020, 0163443720948018. <https://doi.org/10.1177/0163443720948018>.

assessment of the signifying techniques embedded within machine learning and the overarching account of asymptotic labor presented here. One point of reference and inspiration is Miyako Inoue's work, which helps to historicize this argument. Inoue has brilliantly demonstrated how semiotic analysis can be used to trace the ways that specialized labor practices – in her case, stenography in late-nineteenth century Japan, a proto-form of information work, to be sure – can become “*feminized*,” thereby losing their cultural value as skilled intellectual labor.⁸⁸ This mirrors the gendered division of labor latent within the discussion of AI and asymptotic labor, where the decidedly *masculinized* engineers and entrepreneurs are centered within popular discourse, while the essential reproductive labor that undergirds these systems remains hidden from view and systematically undervalued. Julia Elyachar, similarly, has written about the forms of semiotic labor required to create and maintain forms of communications infrastructure, and how these social technologies can “transmit not only language but also all kinds of semiotic meaning and economic value,” which helpfully contributes to our understanding of the role of asymptotic labor in supporting the AI systems described below.⁸⁹ And finally, while it has remained marginal in the multifaceted fields of information studies, several scholars have attempted to unpack the political economy of information work vis-a-vis

⁸⁸ Inoue, Miyako. “Stenography and Ventriloquism in Late Nineteenth Century Japan.” *Language & Communication* - *LANG COMMUN* 31 (July 1, 2011): 181–90. <https://doi.org/10.1016/j.langcom.2011.03.001>. 182.

⁸⁹ Elyachar, Julia. “Phatic Labor, Infrastructure, and the Question of Empowerment in Cairo.” *American Ethnologist* 37, no. 3 (2010): 452–64. <https://doi.org/10.1111/j.1548-1425.2010.01265.x>.

semiotic analysis, including Julian Warner's study of information-retrieval as a process of navigating sign-systems and Jens-Erik Mai's explicit attempt to link the semiotic “index” with the work of cataloging and indexing.⁹⁰

As with my study of CAPTCHA, this article is based primarily on an analysis of the existing technical literature focused on several prominent subtopics in the field of machine learning. However, unlike the preceding chapter, we were less concerned with genealogically mapping the emergence and evolution of a new technological artifact, and rather focused on undertaking a critical close-reading of this source material in order to extract and highlight the semiotic foundations of these techniques, and parlay this into a critical framework that might be fruitful for other scholars interested in studying the relationship between indexical algorithms and labor, and the asymptotic tendencies that subsist at this intersection.

⁹⁰ Warner, Julian. “A Labor Theoretic Approach to Information Retrieval.” *Journal of the American Society for Information Science and Technology* 59, no. 5 (2008): 731–41. <https://doi.org/10.1002/asi.20782>. Mai, Jens-Erik. “Semiotics and Indexing: An Analysis of the Subject Indexing Process.” *Journal of Documentation* 57, no. 5 (January 1, 2001): 591–622. <https://doi.org/10.1108/EUM0000000007095>.

*In investigating history we are not flicking through a series of “stills.” ... Any historical moment is both a result of prior process and an index towards the direction of its future flow.*⁹¹

—E. P. Thompson

The Digital Index: A World Below

We live in the age of indexical artificial intelligence. Data and algorithm combine to direct the flow of commodities and labor power, images and messages, through geographical and institutional space. Avant-garde AI algorithms called neural nets manipulate high volumes of data, forcing us to leave behind the sense that digital computation is abstract, disembodied, heuristic. This technology has become a core element of global capitalist infrastructure, as the parsing of data produced on digital platforms depends on these algorithms. Nets are currently used in targeted advertising, self-driving car technologies, pricing systems, social media, content-delivery platforms like Spotify and Netflix, facial recognition, medical imaging, judicial consulting—the list goes on.⁹² The dangers for economic and racial justice have been widely demonstrated,⁹³ but as we enter the petabyte

⁹¹ E. P. Thompson, “The Poverty of Theory or An Orrery of Errors (1978),” *The Poverty of Theory & Other Essays* (New York, 1978), p. 64.

⁹² See John D. Kelleher, *Deep Learning* (Cambridge, Mass., 2019); hereafter abbreviated *D*.

⁹³ There is an enormous and rapidly growing literature on bias in digital systems. See Julia Angwin and Jeff Larson, “Machine Bias,” *ProPublica*, 23 May 2016, www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing; Timnit Gebru et al., “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” *FAccT ‘21: Proceedings of the 2021 ACM Conference on Fairness,*

era—or perhaps the exabyte or zettabyte era—it seems unlikely that the pace of adoption will slow.⁹⁴

Computers were built to handle symbols assigned arbitrarily to arrangements of hardware with sets of switches operating in rigid syntax. We humans were supposed to be the channel through which the answers to queries flowed back into the world. The scientist at the terminal was supposed to ground the output and assign answers to contexts. But we have plugged our machines into each other and let them take on systemic autonomy, allowing algorithms to set price, determine financial speculation, and find matches between consumer and producer in online-to-offline services like rideshares. Mercedes Bunz has argued that nets “calculate meaning”: the signs that force our attention to our digital devices no longer seem a world apart, a source of consultation, but a world below—

Accountability, and Transparency (Mar. 2021): 610–23; and Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Cambridge, Mass., 2016). For an early perspicacious source, see William H. Dutton and Kenneth Kraemer, “Automating Bias,” *Society* 17, no. 2 (1980): 36–41. See also Paul Baker and Amanda Potts, “‘Why Do White People Have Thin Lips?’ Google and the Perpetuation of Stereotypes via Auto-Complete Search Forms,” *Critical Discourse Studies* 10 (May 2013): 187–204; Safiya Umoja Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York, 2018); and Kate Crawford and Trevor Paglen, “Excavating AI: The Politics of Training Sets for Machine Learning,” 19 Sept. 2019, www.excavating.ai

⁹⁴ See Christopher Rosol, Jürgen Renn, Robert Schlägl et al., “On the Age of Computation in the Epoch of Humankind,” *Nature*, 28 Nov. 2018, www.nature.com/articles/d42473-018-00286-8#ref-CR14. For an important corrective to the idea that the data imaginary emerged only around 2000, see Orit Halpern, *Beautiful Data: A History of Vision and Reason since 1945* (Durham, N.C., 2014).

a semiotic infrastructure.⁹⁵ Digital systems, relying on the neural net, have left the world of mere symbol behind and have begun to ground themselves *here, now, for you*—they are able to *point* to real states of affairs in a sign function known as the index.

We can neither afford to dismiss nor to accept fully the signification of digital systems. As we leave behind the sense that computers merely process symbols unattached from our physical and social realities, we run the risk of believing *too much* that machines really are intelligent. This is why *data-driven* has taken on the ring of necessity in our capitalist institutions. Indexical AI contrasts with the symbolic AI that dominated artificial intelligence research before 2000.⁹⁶ And it uses this powerful referential function, the

⁹⁵ Mercedes Bunz, “The Calculation of Meaning: On the Misunderstanding of New Artificial Intelligence as Culture,” *Culture, Theory and Critique* 60, nos. 3–4 (2019): 264–78. See also M. Beatrice Fazi, “Beyond Human: Deep Learning, Explainability and Representation,” *Theory, Culture & Society*, 27 Nov. 2020, <http://journals.sagepub.com/doi/full/10.1177/0263276420966386>.

⁹⁶ On digital indexicality, see Kris Paulsen, “The Index and the Interface,” *Representations*, no. 122 (Spring 2013): 83–109, and Tom Gunning, “What’s the Point of an Index? Or, Faking Photographs,” in *Still Moving: Between Cinema and Photography*, ed. Karen Beckman and Jean Ma (Durham, N.C., 2008), pp. 39–49. As we were composing this article, we discovered the excellent work of Michael Castelle, who has mentioned in several places that certain aspects of machine learning are indexical. See Jonathan Roberge and Michael Castelle, “Toward an End-to-End Sociology of 21st-Century Machine Learning,” *The Cultural Life of Machine Learning: An Incursion into Critical AI Studies*, ed. Roberge and Castelle (Cham, 2021), p. 9. We agree with Castelle and elaborate the case for an indexical AI here, although we came upon this insight independently of his work (and of each other). On the semiotics of nets more generally, see Paul Kockelman, “The Epistemic and Performative Dynamics of Machine Learning Praxis,” *Signs and Society* 8 (Spring 2020): 1–37. A recent Google *DeepMind* article casts deep learning as “symbolic behavior,” using a semiotic

index, to gain the appearance of truth, the force of necessity. In the semiotic terms that we spell out below, nets indexicalize data to produce judgments about the world—judgments we take to be meaningful and to resemble our own judgments at our own risk. Neural nets disassemble images for recognition and production of new images, pointing us to what appear to be resemblances, images of something other than the pathways of the net itself. They also search for linguistic redundancy, indicating and replicating high-level generic qualities of natural language. Questions about how much the systems *know* about images or language, and through them the world, are missteps and should be abandoned in favor of concrete analysis of these now ubiquitous semiotic operators.

With the spread of these systems, the ability to dissent from their conclusions begins to fade, and the gap between their signification and ours closes. Subtending both the logistics of capital and the inequitable mangling of identity, nets have collapsed the space between infrastructure and representation, the small window of free air between immiseration and humiliation. Indexical AI seems to maintain the vector of history, the feeling of the necessity of capital, reinforcing “capitalist realism,” in Mark Fisher’s phrase.⁹⁷ The sense that there is no alternative is the sense of being borne along on the “future flow” of history, as E. P. Thompson once put it, of simultaneous closure and progress, even if progress is

framework but not focusing on the index (Adam Santoro et al., “Symbolic Behaviour in Artificial Intelligence,” *DeepMind* [2021]: <http://arxiv.org/pdf/2102.03406.pdf>, p. 17).

⁹⁷ See Mark Fisher, *Capitalist Realism: Is There No Alternative?* (Seattle, 2009).

along the darkest timeline. What McKenzie Wark calls “vectoralist” capitalism—or perhaps something “worse” than capitalism—is rooted in the literal indexical vectors of the neural net, which form a material and signifying grid in which we both live and think.⁹⁸

When we accept the material and social reality established by these nets, we engage in what we call the naïve iconic interpretation of AI, an uncritical and ultimately metaphysical belief in the promise of intelligence. Analyzing digital semiotics allows us to peer inside what seem to be “black boxes,” revealing not an icon but an index.⁹⁹ Critique of this new infrastructure must abandon the allied notions that digital systems occupy a merely logical space and that they know better than we do and should guide us.

Net Semiotics

A neural net is a mechanism that learns representations from data. Made up of nodes linked to one another with algorithmic instructions—in simplified form: first multiply, then sum; if less than 0, then delete—nets convene many simple parts to produce complex wholes. A net tasked with classifying canines first undergoes an extended period of training, during which it is exposed to as many relevant examples as is practical. Each image in the training set—an array of pixels plus a semantic label, recast as a single vector

⁹⁸ McKenzie Wark, *Capital Is Dead: Is This Something Worse?* (New York, 2019), pp. 45, 46.

⁹⁹ See Pasquale, *The Black Box Society*.

of information—is fed through the net, multiplied by random weights and a nonlinear function at each node, before finally arriving at the classification stage codified in the net’s output layer. Initially, the net will be highly prone to errors because the parameters that determine how nodes interact with one another begin in a randomized state. By using a technique for algorithmic optimization called stochastic gradient descent, however, the net can assign blame for these errors (or credit for precision) to individual neurons along the pathway from input to output. Over time, this process allows the net to refine the weighted parameters linking nodes gradually and to learn which features in the dataset are significant for detecting, for example, a Samoyed, even distinguishing it from a white wolf that often looks identical to the untrained human eye. The trained net possesses no holistic sense of Samoyedness but rather a complex architecture of indexical pathways that point to Samoyedness by capturing salient relations between features (Figure 3).




Figure 3. A convolutional neural network identifying a white dog.

The semantic range of these algorithms relies on the second function of reference in Charles Sanders Peirce's trichotomy of signs. For Peirce, there are three fundamental modes of designation of an object: the icon, the index, and the symbol. It is common say that the icon signifies by similarity with its object,¹⁰⁰ the index by existentially affecting its object,¹⁰¹ and the symbol by pure convention. The picture in the mind, the portrait and the algebraic equation are icons; the pointing finger, the relative pronoun, and the weathervane are indexes; and the letter, but also the data-point, is a symbol.¹⁰² None of these operates in isolation, and in fact all three are at play in nearly every human use of signs. For example, Peirce identifies an "index of a peculiar kind" that is created by the tokens or instances of symbols in actual use.¹⁰³ Because symbols express only generalities, their instances (words) must really affect other instances or signs in order to ground

¹⁰⁰ See Charles Sanders Peirce, *Principles of Philosophy*, vol. 1 of *The Collected Papers of Charles Sanders Peirce*, ed. Charles Hartshorne, Paul Weiss, and Arthur W. Burks, 8 vols. (Cambridge, Mass., 1932), p. 558.

¹⁰¹ See Peirce, "A Second Trichotomy of Signs," in *Elements of Logic*, p. 248.

¹⁰² There is a rich tradition that extends from Peirce's comment that the photograph is both iconic and indexical (as a trace of light). See especially Roland Barthes, *Camera Lucida: Reflections on Photography*, trans. Richard Howard (New York, 2010), Rosalind Krauss, "Notes on the Index: Seventies Art in America," *October* 3 (Spring 1977): 68–81, and Joel Snyder and Neil Walsh Allen, "Photography, Vision, and Representation," *Critical Inquiry* 2 (Autumn 1975): 143–69. Nets never involve an icon-index relation that lacks a concurrent symbol function, so they never exemplify the indexicality that has been theorized in this particular debate. On digital images see William J. Mitchell, *The Reconfigured Eye: Visual Truth in the Post-Photographic Era* (Cambridge, Mass., 2001).

¹⁰³ Peirce, "A Second Trichotomy of Signs," p. 249.

linguistic discourse. Digital machines, in order to ground their calculations, implement such a peculiar index.

Signs are composite, “indecomposable,” never reducible to a single function.¹⁰⁴ Peirce writes that “a sign is anything which is related to a Second thing, its Object, in respect to a Quality, in such a way as to bring a Third thing, its Interpretant, into relation to the same Object.”¹⁰⁵ The sign is a triadic relation between vehicle, object, and interpretant—it cannot be reduced to less than three.¹⁰⁶ But the sign can *communicate* less than three—the “thirdness” of the sign communicates the “firstness” of a quality or a singularity in the icon and the “secondness” or relation in the index.¹⁰⁷ Icons are “sources of discovery”: “a great distinguishing property of the icon is that by the direct observation of it other truths concerning its object can be discovered than those which suffice to determine its construction” (“I,” p. 279). Peirce gives the example of two photographs that together allow one to draw a map, converting the combined likeness of space into a diagram of territory. “Given a conventional or other general sign of an object,” he writes, “to deduce

¹⁰⁴ Peirce, *Pragmatism and Pragmaticism*, vol. 5 of *The Collected Papers of Charles Sanders Peirce*, p. 469.

¹⁰⁵ Peirce, “Partial Synopsis of a Proposed Work in Logic,” in *Elements of Logic*, vol. 2 of *The Collected Papers of Charles Sanders Peirce*, p. 97.

¹⁰⁶ See Peirce, “Questions Concerning Certain Faculties Claimed for Man,” in *Selected Writings (Values in a Universe of Chance)*, ed. Philip P. Wiener (New York, 1966), pp. 31–32. We think with Peirce in what follows, with the aim of illuminating neural nets, rather than to produce a holistic interpretation of Peirce.

¹⁰⁷ Peirce, “The Icon, Index, and Symbol,” in *Elements of Logic*, p. 276; hereafter abbreviated “I.”

any other truth than that which it explicitly signifies, it is necessary, in all cases, to replace that sign by an icon” (“I,” p. 279). Even algebraic icons—equations—are such a source of discovery because they are made up of signs that communicate qualities, among which we select and compare. Selection is indexical because the index is a sign that communicates a relation (a “second”).¹⁰⁸ Peirce pays special attention to the problem case of logical algebra in the Boolean tradition, which projects a dyadic world of indexes but fails to recognize its own status as a third, or sign-relation.¹⁰⁹ Digital machines, built on an ultimately Boolean algebra that intentionally excludes triadic logic, extend this case and complicate it.¹¹⁰ Indexical AI disproves the exclusion, pushing computing into a semantic domain by means of the peculiar, symbolic index, a third that communicates dyadic

¹⁰⁸ “Every physical force reacts between a pair of particles, either of which may serve as an index of the other. On the other hand, we shall find that every intellectual operation involves a triad of symbols” (Peirce, “The Nature of Symbols,” in *Elements of Logic*, p. 300).

¹⁰⁹ See Peirce, “On Signs and the Categories,” *Reviews, Correspondence and Bibliography*, vol. 8 of *The Collected Papers of Charles Sanders Peirce*, p. 331.

¹¹⁰ The index is often seen as a dual relation, rather than as a sign communicating a second, in a move that takes what the sign communicates for the nature of the sign itself. See Mary Ann Doane, “Indexicality: Trace and Sign: Introduction,” *Differences* 18 (May 2007): 1–6. Another commonly used indexical duality is between presuppositional and creative indexes; see Michael Silverstein, “Shifters, Linguistic Categories and Cultural Description,” in *Meaning in Anthropology*, ed. Keith H. Basso and Henry A. Selby (Albuquerque, 1976), pp. 11–55. Even Silverstein had to admit that this division could not explain uses of “indexical tokens” that bring “into sharp cognitive relief part of the context of speech,” creating out of a presuppositional repertoire something not yet in existence (Silverstein, “Shifters, Linguistic Categories and Cultural Description,” p. 34). See also Constantine V. Nakassis, “Indexicality’s Ambivalent Ground,” *Signs and Society* 6 (Winter 2018): 281–304.

relations among unimaginably many data points. Early nets solved logical problems; current nets make judgments about creditworthiness, jail time, and other nondyadic things. The index in the neural net substitutes for and manipulates the other sign functions, heightening the sense that net output is meaningful, trustworthy, robust. Terrence Deacon argues that “only indexical relationships provide information,” while “iconic relationships” can be used to acquire information, and symbols constitute “relationships between forms of information.”¹¹¹ The sheer complexity of this semiotic function, combined with the volume of data and amount of processing, causes us to lose sight of this information-pushing function, this indexical information.

Since Alan Turing threw down the gauntlet for machine intelligence, the question of imitation has been a backbone of AI research. Turing famously claimed that the way a machine could prove its intelligence would be by becoming indistinguishable from a human

¹¹¹ The index is often seen as a dual relation, rather than as a sign communicating a second, in a move that takes what the sign communicates for the nature of the sign itself. See Mary Ann Doane, “Indexicality: Trace and Sign: Introduction,” *Differences* 18 (May 2007): 1–6. Another commonly used indexical duality is between presuppositional and creative indexes; see Michael Silverstein, “Shifters, Linguistic Categories and Cultural Description,” in *Meaning in Anthropology*, ed. Keith H. Basso and Henry A. Selby (Albuquerque, 1976), pp. 11–55. Even Silverstein had to admit that this division could not explain uses of “indexical tokens” that bring “into sharp cognitive relief part of the context of speech,” creating out of a presuppositional repertoire something not yet in existence (Silverstein, “Shifters, Linguistic Categories and Cultural Description,” p. 34). See also Constantine V. Nakassis, “Indexicality’s Ambivalent Ground,” *Signs and Society* 6 (Winter 2018): 281–304.

interlocutor in the imitation game. The theory of intelligence has always had this iconic root—because we do not know the mechanics of human intelligence (at least not in interaction with other humans), *resemblance* is the crucial feature of any AI.¹¹² One can try to solve this problem of likeness by approximating the expression of intelligence or by imitating the underlying hardware, the brain. Both approaches take the surface feature of the icon—likeness—as the crucial element of AI. Machines must either be intelligent from similar physical resources to ours, or they must at least express themselves in signs that could only come from an intelligence like ours.

We call this implicitly iconic theory of intelligence the naïve iconic interpretation because the comparison rests on a resemblance unanchored in the expression of a *first* or quality.¹¹³ The comparison of human and machine intelligence is slippery. The implicit icon (the quality intelligence) has no known domain—we are supposed to know it when we see it. And of course we do, which is why AI proceeds by provocative exhibitions—as the automata tradition did from antiquity on¹¹⁴—from Ray Kurzweil’s machine-composed

¹¹² See A. M. Turing, “Computing Machinery and Intelligence,” *Mind* 59 (Oct. 1950): 433–60.

¹¹³ We can distinguish this from a non-naïve iconic interpretation like that of Nils Nilsson, who divides AI representations into “feature” maps (descriptive, incomplete) and “iconic” (simulative, complete) (Nils J. Nilsson, *Artificial Intelligence: A New Synthesis* [San Francisco, 2003], pp. 71, 74). Both forms tend to rely, we suggest, on indexical AI.

¹¹⁴ See Jessica Riskin, *The Restless Clock: A History of the Centuries-Long Argument over What Makes Living Things Tick* (Chicago, 2016).

piano piece in 1965 to Lee Sedol’s defeat at the hands of the neural net AlphaGo in 2015.¹¹⁵ These exhibitions always beg the question. The neural net *might* be like human intelligence in many different ways, but the naïve iconic interpretation is a metaphysical misstep. Neural nets operate indexically, as a quick glance at a canonical diagram of one suggests (Figure 4).¹¹⁶




Figure 4. Diagram of a neural net.

¹¹⁵ See Sidewinder77, *Ray Kurzweil on “I’ve Got a Secret,”* YouTube, 18 Apr. 2007, www.youtube.com/watch?v=X4Neivqp2K4&ab_channel=Sidewinder77. See also George Johnson, “To Test a Powerful Computer, Play an Ancient Game,” *New York Times*, 29 July 1997, www.nytimes.com/1997/07/29/science/to-test-a-powerful-computer-play-an-ancient-game.html.

¹¹⁶ See Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, “Deep Learning,” *Nature* 521, 28 May 2015, pp. 436–44.

Data inputs are treated as indexes, fed forward and then backpropagated through the net.¹¹⁷ Arrows—as pure an index as a diagram can feature—are ubiquitous in these visualizations. Nets are often characterized as pattern detectors, but they do not see patterns as a whole. At a depth of five to twenty hidden layers, the net can indeed distinguish Samoyeds from white wolves with striking accuracy, as Yann LeCun, Yoshua Bengio, and Geoffrey Hinton report.¹¹⁸ But the comparison is limited to pointing. The job of the net is to form capillary connections between informational units and apply these connections after training. Whatever surplus in apparent intelligence it gains is given in the array of those units. When Bengio writes that “Deep learning allows the computer to build complex concepts out of simpler concepts,” with “each concept defined in relation to simpler concepts” and “more abstract representations computed in terms of less abstract ones,” the entire question of intelligence is begged.¹¹⁹ The question of whether a net possesses a concept is being posed rather too early. The net *finds its way* to a pattern—Samoyed, including “not-white-wolf”—through iterative vector multiplication, creating a pathway that generalizes from object to class. Even the core metaphor, *learning*, renders an iconic interpretation of an indexical process. The metaphor suggests that this learning

¹¹⁷ To avoid confusion: signs, including the index, are not necessarily interpreted or interpretable by a human mind, meaning that a strict separation between the mathematical operations of hardware and the signifying operations of an algorithm falls apart under the semiotic microscope; see Peirce, *Principles of Philosophy*, p. 339 and *Elements of Logic*.

¹¹⁸ See LeCun, Bengio, and Hinton, “Deep Learning,” p. 438.

¹¹⁹ Ian Goodfellow, Bengio, and Aaron Courville, *Deep Learning* (Cambridge, Mass., 2016), pp. 5, 8.

is like ours—or like intelligence as such—but does not specify how. The comparison buries the specificity that would help us understand these algorithms. Our proposal is that a version of Peirce’s peculiar index is at work here—that we have less an artificial intelligence than an automated index generator. Images and language, among other types of data, are disassembled, tokenized, and their iconic contents rearranged indexically. The output can be reintegrated into our collective knowledge, even given the status of an icon—everything has qualities, after all. But we should not accept the force of net-produced conclusions as though they were iconically grounded. The metaphor *learning* papers over the semiotic process, contributing to adoption but not critical analysis.¹²⁰

This is not to say that the net does not generate hypotheses. When Spotify’s algorithm expresses the determinate judgment “this is a grunge song,” we should strictly state this as “if grunge ranges over these songs, this specific song can be included.” Judgments of this sort are called “abduction.” François Chollet captures this ability of nets nicely.¹²¹ Rather than applying a general law to phenomena, or deriving a tentative law from phenomena, abductive reasoning is hypothetical; if there were a general law that included

¹²⁰ See Sam P. Kellogg, “The Mountain in the Machine: Optimization and the Landscapes of Machine Learning,” *Culture Machine* 120 (2021).

¹²¹ See François Chollet, *Deep Learning with Python* (Shelter Island, N.Y., 2018), p. 5. Leo Breiman famously called this sort of approach “algorithmic modeling.” Breiman, Leo. “Statistical Modeling: The Two Cultures (with Comments and a Rejoinder by the Author).” *Statistical Science* 16, no. 3 (August 2001): 199–231. <https://doi.org/10.1214/ss/1009213726>.

these examples, *then* this particular item would fall under it (Figure 5).¹²² The net finds the function rather than applying it. Luciana Parisi has argued that neural nets display just this kind of judgment, rendering them “transcendental[ly] instrumental”; they are tools that possess a sort of judgmental autonomy.¹²³ The icon drives abductive reasoning because we must identify and differentiate qualities by direct comparison to arrive at a hypothesis. But a neural net does not quite reason iconically, as we will argue.¹²⁴ It generates hypotheses in the form of indexical pathways that it traces between the tokens it gets as input. Pixels, for example, combine to make an image, which is tempting to call an icon (and may be interpreted as one). But the action of the net occurs on a disassembled icon, training on and producing an indexical pattern of tokens that can *choose* to construe as an icon once again—or not. In this respect, even the image fed through the neural net is more like a literary icon, a quality communicated in hierarchically arranged symbols, than an image understood as resembling its referent.¹²⁵ The naïve iconic interpretation of

¹²² See Peirce, *Elements of Logic*, p. 619.

¹²³ Luciana Parisi, “Reprogramming Decisionism,” *e-Flux* 85 (Oct. 2017), www.e-flux.com/journal/85/155472/reprogramming-decisionism/.

¹²⁴ In a strict sense, neural nets *are* icons of a specific sort because they are extended equations—logical icons with complex nonlogical content. But neither humans nor nets themselves can make *sense* of these equations, as the layers are many and hidden. They are God’s-eye icons. And nets do not possess the semantic range of human icons, partly because they accept disassembled human icons as their inputs and are dependent on them. They are unable to integrate perceptual, logical, and verbal icons—the full semantic stack of human icon use.

¹²⁵ See note above.

learning systems rests on this confusion. It takes the image as icon, the ability to identify images as iconic reasoning. Neural nets, however, are indexical all the way through.




Figure 5. Machine Learning outputs rules.

Indexical AI

At the opposite extreme of the iconic interpretation of neural nets is the paradigm known as symbolic AI, which dominated AI research from its inception in the 1950s to the turn of the millennium. During this period, the paradigm that now encompasses deep learning went under many names—machine learning, parallel distributed processing, and connectionism.¹²⁶ All have shared the conception of the neural net, and all to some extent

¹²⁶ An excellent historical overview is given in Cameron Buckner and James Garson, *Connectionism and Post-Connectionist Models* (New York, 2018). See also Cardon, Dominique, Jean-Philippe Cointet, and Antoine Mazières. “Neurons Spike Back (La Revanche Des Neurones: L’invention Des Machines Inductives et La Controverse de l’intelligence Artificielle).” *Réseaux* n° 211, no. 5 (2018): 173. <https://doi.org/10.3917/res.211.0173>.

are “inspired by the actual biology of the brain,”¹²⁷ with neuroscientists and AI researchers working interdisciplinarily, sharing an epistemology opposed to a “symbolic” construal of intelligence.¹²⁸

In their 1975 Turing Award speech, “Computer Science as Empirical Inquiry”—the canonical explanation of the symbolic AI program—Herbert Simon and Alan Newell write that a “physical symbol system consists of a set of entities, called symbols, which are physical patterns that can occur as components of another type of entity called an expression (or symbol structure).”¹²⁹ These patterns are able to designate or interpret objects belonging to the system’s domain. The project of building digital computers in the first place was partly the construction of artificial languages in which hardware could operate, to carry out operations on command, answering queries in numbers or language. The symbol system is unequivocal and arbitrary in its designation of values and hierarchically coded, in keeping with Peirce’s definition of the symbol, or Ferdinand de Saussure’s famous “arbitrariness of the sign.”¹³⁰ Preloaded with meaning by imputation,

¹²⁷ Terrence J. Sejnowski, *The Deep Learning Revolution* (Cambridge, Mass., 2018), p. 32.

¹²⁸ Patricia Smith Churchland and Sejnowski, *The Computational Brain* (Cambridge, Mass., 2017), p. 423.

¹²⁹ Allen Newell and Herbert A. Simon, “Computer Science as Empirical Inquiry: Symbols and Search,” *Communications of the ACM* 19 (Mar. 1976): 116.

¹³⁰ Ferdinand de Saussure, *Course in General Linguistics*, trans. Wade Baskin, ed. Perry Meisel and Haun Saussy (New York, 2011), p. 131.

these systems sever the connection between semantic unit and reference, gaining power at the syntactic level purchased at the price of worldly knowledge.

The connectionists thought the symbolic picture was incomplete and sometimes spoke of a “subsymbolic paradigm.”¹³¹ Symbols have semantic content arbitrarily but unequivocally associated with them. But what if intelligence is made up not of symbols but of connections between semantic units? The suspicion that the brain’s neurons generate mental content through weighted links led to this alternate proposal.¹³² What is underneath the symbols in the “subsymbolic paradigm” is also *between* the symbols—links, connections: in a word, indexes. Nonsymbolic AI, then, is always a form of indexical AI.¹³³

¹³¹ Rumelhart, Hinton, and Williams, “Learning Internal Representations by Error Propagation,” in *Psychological and Biological Models*, in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, ed. Rumelhart and James L. McClelland, 2 vols. (Cambridge, Mass., 1999), 2:195. See also Paul Smolensky, “On the Proper Treatment of Connectionism,” *Behavioral and Brain Sciences* 11, no. 1 (1988): 1–23.

¹³² See Cameron Buckner and James Garson, “Connectionism,” *Stanford Encyclopedia of Philosophy Archive*, 16 Aug. 2019, <http://plato.stanford.edu/archives/fall2019/entries/connectionism>.

¹³³ The debate between symbolists and connectionists was always about where semantic content comes from, with symbolists like Jerry Fodor arguing that semantic content comes in systemic “clumps” that could never emerge from a series of links between neurons or vectors (Jerry Fodor and Brian McLaughlin, “Connectionism and the Problem of Systematicity: Why Smolensky’s Solution Doesn’t Work,” *Cognition* 35 [May 1990]: 184). Some connectionists try to include some symbolic processing, while “radical” connectionists believe in pure emergence of the symbolic function (Buckner and Garson, “Connectionism”).

There is a critical tradition in the history of AI that has sometimes conceived of intelligence as indexical. Its main proponent is Philip Agre, who places himself in the phenomenological tradition of AI critique that began with Hubert Dreyfus’s Heideggerean attack on symbolic AI in the late 1960s.¹³⁴ Arguing that “indexicality has been almost entirely absent from AI research,” Agre proposed a “deictic ontology”—achieving contextual awareness by physical indexes—which “can be defined only in *indexical* and *functional* terms, that is, in relation to an agent’s spatial location, social position, or current or typical goals or projects.”¹³⁵ This tradition conceives of intelligence as indexical and so makes indexicality an aspiration of an alternate AI. But this aspiration has blinded the tradition, which still seeks to realize “intelligence,” to the indexical effects of actually existing AI. Take, for example, Brian Cantwell Smith’s recent argument that the crucial concept for deep learning is *context*:

The issue is not merely one of having a computer system use context-sensitive structures and symbols in appropriate ways, such as indexicals and perspectival descriptions (analogous of “today,” or “the medium in this drive,” etc.), but of

¹³⁴ See Leif Weatherby, “Intermittent Legitimacy: Hans Blumenberg and Artificial Intelligence,” *New German Critique* 48 (Feb. 2022).

¹³⁵ Philip Agre, *Computation and Human Experience* (New York, 1997), pp. 241, 243. See also Brian Cantwell Smith, *On the Origin of Objects* (Cambridge, Mass., 1998), p. 168 and *The Promise of Artificial Intelligence: Reckoning and Judgment* (Cambridge, Mass., 2019).

configuring a system’s [deliberation] to be appropriate to the wider situation at hand beyond what is immediately represented, either explicitly or implicitly.¹³⁶

The ability to differentiate context in this way is iconically grounded. So long as we aim at the creation of intelligence, we remain in the discourse of the icon, which allows for flexible differentiation of the sort envisioned here. But deep learning systems have managed to create indexes without the type of intentional awareness of environment that Smith or Agre imagine. Cameron Buckner has suggested their success is due to a kind of “transformational abstraction,” of the kind that John Locke imagined, in which the net can recognize a class from examples and generate exemplars based on knowledge of the class.¹³⁷ This view is a non-naïve version of the iconic interpretation, as it suggests that nets compare qualities and reasoning in terms of identity and difference—achieving the connection of perception and reference that Agre envisioned in deixis. We argue instead that nets obviate the need for perceptual icons or deixis by means of an algorithmic index with both semantic power and semiotic limitation.

¹³⁶ Smith, *The Promise of Artificial Intelligence*, p. 138. This notion, as he notes, is closely aligned with the DARPA’s notion of a “third-wave” AI (p. 138).

¹³⁷ See Buckner, “Empiricism without Magic: Transformational Abstraction in Deep Convolutional Neural Networks,” *Synthese* 195 (Sept. 2018): 5339–72.

The Shape of a Neural Net

The net is a complex function with a concrete shape. In training, the net gains the ability to point at features, and composite sets of features add up to objects. Because the output is then either a determinate judgment (“this is a Samoyed, not a white wolf”) or an image or text that resembles human pictorial or linguistic content, we are predisposed to think of that output as iconic. But the net does not operate iconically; it treats the iconic aspects of its input indexically. Its rhetorical power and semiotic limitation come from its shape.

What has become the neural net was first proposed as a “nervous net” by cybernetician Warren McCulloch and his prodigy student Walter Pitts.¹³⁸ They sketched diagrams of formalized neurons that would allow either a brain or a machine (not yet a computer in 1943) to perform Boolean operations.¹³⁹ McCulloch and Pitts thought that a weight-conditional set of neurons, if complex enough, might be able to achieve true knowledge of the world (Figure 6).¹⁴⁰

¹³⁸ Tara H. Abraham, *Rebel Genius: Warren S. McCulloch’s Transdisciplinary Life in Science* (Cambridge, Mass., 2016), p. 113.

¹³⁹ See *ibid.*

¹⁴⁰ See Warren S. McCulloch and Walter Pitts, “A Logical Calculus of the Ideas Immanent in Nervous Activity,” *Bulletin of Mathematical Biophysics* 5 (Dec. 1943): 115–33. See also Weatherby, “Digital Metaphysics,” *Hedgehog Review* (Spring 2018), <http://hedgehogreview.com/issues/the-human-and-the-digital/articles/digital-metaphysics>.




Figure 6. Logical circuits in a “nervous net” (McCulloch and Pitts, 1943).

The diagram (c) shows an AND function: if (1) and (2) are both tripped across the synapse, (3) fires, but not if only one or the other does. Many early machine learning systems were built on this insight, including Frank Rosenblatt’s perceptron.¹⁴¹ But these machines were unable to perform an XOR function, in which the neuron fires if one or the other of its inputs trips, but not both. It was this problem that proved the efficacy of *deep* nets, nets with hidden layers, meaning the calculations at these layers are inaccessible to the input and output stations.

A single hidden unit turned out to be enough to solve the XOR problem. As David Rumelhart, Geoffrey Hinton, and Ronald Williams were able to show, a hidden unit allowed the net to learn weights that rarely failed to execute XOR.¹⁴² The middle layer

¹⁴¹ See Marvin Minsky and Seymour Papert, *Perceptrons: An Introduction to Computational Geometry* (Cambridge, Mass., 1988).

¹⁴² See Rumelhart, Hinton, and Williams, “Learning Internal Representations by Error Propagation.”

creates a step between input and output that gives the calculations shape—this is why nets are described as having architecture. This early diagram visualizes the flow of calculations labeled with the learned parameters to solve XOR (Figure 7). If we look at the truth tables for inclusive or (OR) and XOR, we can see in semiotic terms why a hidden layer is needed to solve this equation (Figure 8).




FIGURE 2. A simple XOR network with one hidden unit. See text for explanation.

Figure 7. Solving for XOR with one hidden layer.

OR

input	input	output
0	0	0
0	1	1
1	0	1
1	1	1

XOR

input	input	output
0	0	0
0	1	1
1	0	1
1	1	0

Figure 8. Logic tables for OR and XOR functions.

It is tempting to think of these tables as pure binary symbolic expressions, as they contain only logical tokens with binary values. But equations—even when they are not laid out in tables—are icons displaying the reasoning they denote (Figure 9).¹⁴³ They are signs that communicate a “first,” in this case the quality of a judgment or style of reasoning, and therefore “sources of discovery.” But while OR operates on one binary logic (the presence of 1 as input means 1 as output), XOR shifts to another, in which two trues become a false, basing the value 1 on the presence of identical or different inputs. This requires an extra step in reasoning, a pathway through the icon that dissociates binary arithmetic from the query. It is this extra step that the hidden layer allows for. As

¹⁴³ See Peirce, *Elements of Logic*, 2:279.

Rumelhart and his colleagues noted, full isolation of input from output required even more layers and led to even more general function finding.




FIGURE 4. A simple architecture for solving XOR with two hidden units and no direct connections from input to output.

Figure 9. XOR with hidden layers.

If XOR is an algebraic icon, then we can see that the net does not possess that icon as input. Instead, it establishes an indexical pattern that generalizes to instances of the function. In the case of a Boolean function, if the output is correct, it must be a matter of indifference whether the algorithmic architecture is iconic or indexical. But this indifference will not survive the introduction of petabytes of social data. Early nets were able to solve logical problems without having the icon as a source of discovery, a way of ensuring that truths learned from the data actually held. Logic itself made those assurances in problems like these. But if the source of discovery for data about

creditworthiness, racial identity, or market fluctuations is not iconic, no a priori set of rules comes to our aid in reading the indexical results.

When humans input icons into deep learning systems, asking them to produce real answers to real questions, these systems indexicalize the icons they are asked to handle. After the XOR problem, Rumelhart, Hinton, and Williams introduce symmetry as a problem that only hidden layers can solve. Symmetry cannot be detected in the comparison of sums: “an individual input unit, considered alone, provides no evidence about the symmetry or non-symmetry of the whole input vector, so simply adding up the evidence from the individual input units is insufficient.”¹⁴⁴ Finding these particular balance points is the work of the “intermediate layers.”¹⁴⁵

Symmetry is the concept of internal iconicity. A symmetrical object resembles itself across an axis. Even more obviously than in the case of the XOR function, the detection of symmetry requires the net to trace two lines of equivalency. When you or I look at the figure, we can detect symmetry by scanning the parallel numbers—a slightly different operation but with the same result, as when I see an apple as symmetrical (Figure 10). The net, however, cannot rely on the comparison of arbitrary points against the whole to

¹⁴⁴ Rumelhart, Hinton, and Williams, “Learning Representations by Back-Propagating Errors,” *Nature* 323, 9 Oct. 1986, p. 535.

¹⁴⁵ *Ibid.*, p. 533.

check symmetry. At no point in the forward or backward passes can the net compare the whole shape to an individual point (this is the meaning of hidden layers). Looked at from the outside, the pathways that the net establishes make up an icon. *But no one—nothing—looks at them from the outside.* The reason the icon is a source of the discovery of truth not contained in the rules of its construction is that a human interpretant of an icon can compare it to its referent. But the source of discovery in the neural net is not the icon as such. Any truth the net discovers must be a combination of the symbolic input and the indexical pathways that the net establishes—the source of discovery is the pattern through the layers: indexical tokens arranged by backpropagation.¹⁴⁶ This peculiar sign function is the digital index.

¹⁴⁶ Rumelhart, Hinton, and Williams speak of solutions the net discovers that they had not previously seen or considered: “We frequently discover these more elegant solutions by giving the system more hidden units than it needs and observing that it does not make use of some of those provided. Some analysis of the actual solutions discovered often leads us to the discovery of a better solution involving fewer hidden units” (Rumelhart, Hinton and Williams, “Learning Internal Representations by Error Propagation,” p. 341).




Fig. 1 A network that has learned to detect mirror symmetry in the input vector. The numbers on the arcs are weights and the numbers inside the nodes are biases. The learning required 1,425 sweeps through the set of 64 possible input vectors, with the weights being adjusted on the basis of the accumulated gradient after each sweep. The values of the parameters in equation (9) were $\epsilon = 0.1$ and $\alpha = 0.9$. The initial weights were random and were uniformly distributed between -0.3 and 0.3 . The key property of this solution is that for a given hidden unit, weights that are symmetric about the middle of the input vector are equal in magnitude and opposite in sign. So if a symmetrical pattern is presented, both hidden units will receive a net input of 0 from the input units, and, because the hidden units have a negative bias, both will be off. In this case the output unit, having a positive bias, will be on. Note that the weights on each side of the midpoint are in the ratio 1:2:4. This ensures that each of the eight patterns that can occur above the midpoint sends a unique activation sum to each hidden unit, so the only pattern below the midpoint that can exactly balance this sum is the symmetrical one. For all non-symmetrical patterns, both hidden units will receive non-zero activations from the input units. The two hidden units have identical patterns of weights but with opposite signs, so for every non-symmetric pattern one hidden unit will come on and suppress the output unit.

Figure 10. A neural net solving for symmetry.

Once there are hidden layers the system cannot learn merely by comparing input and output. This is because the output alone gives no sense of where the algorithm's weights should be adjusted to optimize the outcome, known as the "credit (or blame) assignment problem" (D, p. 186). The technique that solved this problem—which is regarded as the crucial step to the spread of nets around the turn of the millennium—is the

backpropagation algorithm.¹⁴⁷ As John Kelleher points out, the entire process of correcting error in a net is often referred to as “backpropagation,” a two-step process—measuring the error, updating the weights (*D*, p. 209). Once the net has performed the “forward pass” (multiplication and sums of the layers), there follows the “backward pass,” which assigns credit and blame (an “error value,” ∂) to each node.¹⁴⁸ Although these values are summed, the layers are not visible to the output or the input, with the result that the net takes the shape of the indexical patterns it establishes between the layers in the passes.¹⁴⁹ The pattern includes the input, output, and individual points, but the unity of this pattern—the shape of a neural net—is not available to the net and in a sense also does not exist for *us*, as we cannot make sense of its extremely high volume of data and processing at each neuron. The pattern exists only indexically, as the throughput of the algorithm. In distinguishing *dog* from *wolf*, the net is not comparing dogs and wolves. It is pointing by means of a distributed series of indexical tokens at the label *Samoyed*, and it does this by means of the digital, or algorithmic, index.

Indexing Images

The digital index is a semiotic operation that drives net-based image processing, which might otherwise appear to be an iconic operation. The images in Figure 11 were produced

¹⁴⁷ See Goodfellow, Bengio, and Courville, *Deep Learning*, p. 225.

¹⁴⁸ *Ibid.*, pp. 214, 113.

¹⁴⁹ See *Ibid.*, p. 221.

by a generative adversarial neural net (GAN) designed by engineers at NVIDIA, in which two nets are placed in competition with one another, one finding flaws in the other's output (Figure 11). The faces in this image do not belong to any humans. This sort of output is hard to resist taking as a source of discovery—surely even if we avoid being tricked by these images, we must think they have captured some quality like faciality. But the icon face is produced, in this case as in all nets, indexically.



Figure 2. Uncurated set of images produced by our style-based generator (config F) with the FFHQ dataset. Here we used a variation of the truncation trick [42, 5, 34] with $\psi = 0.7$ for resolutions $4^2 - 32^2$. Please see the accompanying video for more results.

Figure 11. Images produced using a generative adversarial neural network.

This GAN relies on convolutional neural nets (CNN), which have proven particularly well-equipped for classification and regression tasks involving the sensory domains of vision

and voice, recognizing the faces in your Facebook uploads or parsing your commands to Siri. These use cases tend to involve inputs organized into an array structure, whether two-dimensional bitmaps or one-dimensional time-sequence data. Like all contemporary neural networks, CNNs operate with backpropagation but draw much of their interpretive power from the interplay of distinctive architectural elements. This architecture has proven to be generative precisely inasmuch as it is capable of achieving feats of emergence that are ostensibly incompatible with the digital medium.¹⁵⁰ CNNs demonstrate that digital information is hardly limited to the symbolic register, and computers are anything but mere symbol manipulators. Yet they also are not in the business of processing icons, however much we may tend to think of image and icon as adjacent. CNNs indexicalize images, including their iconic aspects, and either class them or generate proximate images by means of the digital index.

The architectural elements of the CNN responsible for the indexicalization of images were laid out in 1980 by Kunihiro Fukushima.¹⁵¹ He called his framework the “neocognitron,” and it consisted of a multilayered structure of “simple” and “complex” cells.¹⁵² Inspired by the landmark studies undertaken by neurophysiologists David Hubel and Torsten Wiesel

¹⁵⁰ See Brian Massumi, “On the Superiority of the Analog,” *Parables for the Virtual: Movement, Affect, Sensation* (Durham, N.C., 2002), pp. 133–43.

¹⁵¹ See Kunihiro Fukushima, “Neocognitron: A Self-Organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position,” *Biological Cybernetics* 36 (Apr. 1980): 193–202.

¹⁵² *Ibid.*, p. 193.

that discovered a hierarchically structured neural network in the feline visual cortex, the neocognitron architecture consisted of alternating layers of S-cells and C-cells (Figure 12).¹⁵³ S-cells distinguish excitatory from inhibitory regions, working like a flashlight scanning a dark interior, illuminating discrete parts of the uncertain whole. The S-cell performs something like the gestalt operation of interleaving foreground and background, bringing the two into relation with one another. The C-cell takes as its input an amalgamation of the bounded regions captured by the preceding S-cells and congeals them into a higher-order output. Because a single C-cell effectively transforms the features captured by different S-cells, the simple-to-complex data exchange between these layers renders an output that is comparatively agnostic to the spatial (or temporal) contiguities that the S-cell is designed to respond to. S-cells want to know *where* and *when* certain features exist, while C-cells only need to know *that* these features exist. The neocognitron adumbrated the crucial feature of contemporary CNNs: S-cells burned indexes through the data, while C-cells pooled those indexes, creating an indexical sum to identify an object in the image.

¹⁵³ See D. H. Hubel and T. N. Wiesel, "Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex," *The Journal of Physiology* 160 (Jan. 1962): 106–54.




Fig. 5. An example of the interconnections between cells and the response of the cells after completion of self-organization

Figure 12. Diagram of Fukushima's neocognitron, consisting of alternating layers of S-cells and C-cells.

Beginning in the late 1980s, LeCun, working with a rotating slate of collaborators, integrated the promising new technique of backpropagation into the convolutional architecture,¹⁵⁴ developing LeNet to recognize handwritten letters.¹⁵⁵ LeNet's

¹⁵⁴ See Lecun, "A Theoretical Framework for Back-Propagation," *Proceedings of the 1988 Connectionist Models Summer School, CMU* (Pittsburg, 1988), pp. 21–28.

¹⁵⁵ See LeCun et al., "Backpropagation Applied to Handwritten Zip Code Recognition," *Neural Computation* 1, no. 4 (1989): 541–51 and "Handwritten Digit Recognition with a Back-Propagation Network," *Advances in Neural Information Processing Systems (NIPS 1989)* (Denver, Colo., 1990),

convolutional layers produce indexical feature maps, and pooling (or subsampling) layers reassemble these features into potential icons. Finally, for the net to render a judgment interpretable to humans, a fully connected output layer converts the input from the final pooling layer into another index, flattening the multidimensional array into something resembling a traditional look-up table (not dissimilar from a book's index). This duo of convolution and pooling enables the net to realize qualitative novelty out of a quantitative substrate. CNNs, ultimately, are two-part navigational devices: convolutional layers carve new informational pathways through uncharted latent space, and pooling layers groom these pathways to ensure easier transit.

Consider a digital image: twenty-eight pixels tall, twenty-eight pixels wide (Figure 13). This array of 784 pixels presents to the human eye a grayscale raster depicting a handwritten digit. It is one of 9,298 images provided to LeCun's team by the US Postal Service, culled and normalized from a database of zip codes. Because the image is black-and-white, the input layer has the dimensions $28 \times 28 \times 1$, and the output layer is 10×1 , each neuron providing a weighted probability that the input contains one of the ten possible digits.¹⁵⁶ Four layers in LeNet play alternating roles: H1 and H3 are convolutional

<http://nyuscholars.nyu.edu/en/publications/handwritten-digit-recognition-with-a-back-propagation-network-2>.

¹⁵⁶ If the image were color, there would be three channels of information per pixel position—associated with red, green, and blue values—making for a higher-dimensional input shape with the volume $28 \times 28 \times 3$

layers, while H2 and H4 are subsampling layers. Convolutional layers, in essence, add depth: H1 puffs up the flat, two-dimensional input array into a stack of four arrays, each dedicated to a specific feature detection task. The subsampling operation performed by H2, in turn, retains the fourfold depth model produced by H1 but reduces the breadth of each layer by half. H3 and H4 continue this process: the former expands the input received from H2, tripling its depth from four feature-detection planes to twelve, and sends this narrow-but-deep data to H4, where it is narrowed by half once again.¹⁵⁷

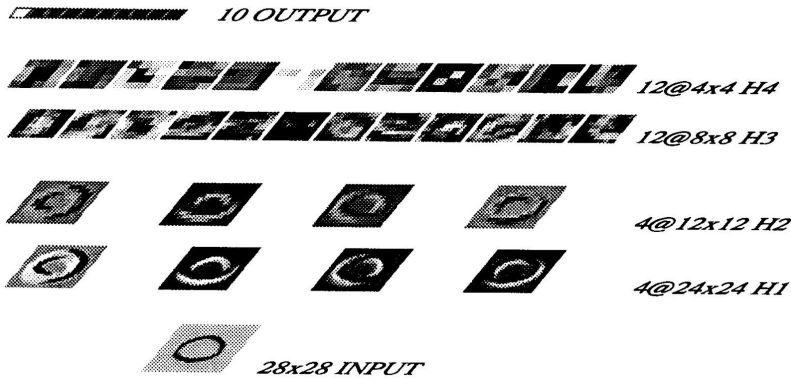


Figure 4: Network Architecture with 5 layers of fully-adaptive connections.

Figure 13. Feature maps produced by each layer of LeNet.

shape. Adding dimensions has caused compute requirements to explode or go parabolic; see “AI and Compute,” *OpenAI*, 16 May 2018, <http://openai.com/blog/ai-and-compute/>.

¹⁵⁷ See figure 11.

Convolution is a mathematical operation that produces a third function out of two input functions. The first function is the input, the second function is the kernel (a matrix that passes over the data image, checking for the desired features), and the third function is the feature map. The feature map is a function that represents the influence of the kernel on the input, and it is here that we see the net indexicalize the image, in a process that resembles scanning. H1 produces a feature map comprising new signifying elements that represent the influence of convolving the values in a 5×5 kernel over the values in each possible 5×5 region of the input. The indexical tensor that this operation produces has the dimensions $24 \times 24 \times 4$, the latter value indicating the number of unique filters that have convolved the input data. The four feature maps in LeNet's H1 layer depict what the net has determined to be salient for recognizing the digit 0. These features serve as heuristics for discovering granular primitives like edges, curves, and contours.

On first glance, H2 seems to do similar work to H1. But while both operate on local patches of their input, with H2 employing a smaller kernel than its predecessor, the two layers differ markedly in how they treat this data. Unlike the convolutional filter, which moves across the input in overlapping strides, the receptive fields in the subsampling layer do not overlap. Likewise, they perform a more elementary bit of math than H1, by simply taking a local average of the values in each receptive field. This averaging function is inelegant compared to the dynamic processing of convolution, but its aim is to preserve as much of the indexical information elicited by the preceding layer, while maximizing the

net's efficiency by reducing the overall breadth. The tensor produced at this stage retains the layers it received as input from the convolutional layer, but the resolution of each layer has been halved, producing a new $12 \times 12 \times 4$ mapping of the data's features. The representations produced by pooling trade resolution for resilience, ensuring like features can be detected as like, despite the innumerable contingencies in real-world input data. Pooling endows the net with translational invariance allowing it to redeploy the features learned in the convolutional stage with more accuracy and efficiency. Pooling kernels essentially force the net to treat inputs with slight variance as like one another. Unlike the indexicalization performed by the convolutional layer, pooling layers like H2 determine which of the differences in the input data can be ignored, and which differences are salient. By instructing the net *not* to notice certain differences, H2 exercises a thresholding operation to produce a map of potential icons, which will be shuttled onward to H3 for further indexicalization. Recall Deacon's argument, which captures this process: "Only indexical relationships directly provide information ... iconic relationships [present] the possibility of being used to acquire information."¹⁵⁸ The CNN architecture is designed to maximize potential iconicity by stacking convolutional and pooling layers in the indexical

¹⁵⁸ Deacon, "Shannon – Boltzmann – Darwin," p. 191. Elsewhere, Deacon stresses that "similarity does not cause iconicity," illustrating that iconicity is always dependent upon the conditions of interpretation (Deacon, *The Symbolic Species: The Co-Evolution of Language and the Brain* [New York, 1997], p. 71). In the case of CNNs, these conditions are codified by the parameters of the pooling layer.

pathway (Figure 14).¹⁵⁹ H3 seizes on this iconic potentiality and produces a third feature map with even greater depth. H4 retains this newfound depth, while once again halving the breadth producing a final tensor with the dimensions $4 \times 4 \times 12$. Each of the 192 neurons in this final pooling layer are connected to the output layer of LeNet, which bears ten total neurons, one for each possible digit. Unlike the generative indexing performed by convolutional layers, this fully connected layer compresses the feature map of H4 into a rigid table of reference, yoked to a set of predetermined tokens—another index.




Figure 1: Architecture of LeNet 1. Each plane represents a feature map, i.e. a set of units whose weights are constrained to be identical. Input images are sized to fit in a 16 x 16 pixel field, but enough blank pixels are added around the border of this field to avoid edge effects in the convolution calculations.

Figure 14. Architecture of the first implementation of LeNet, illustrating the relationship between convolutional and pooling layers.

¹⁵⁹ The linguistic anthropologists Susan Gal and Judith T. Irvine have written at length on this process of treating indices as icons, deeming it a critical aspect of all meaning-making systems. They name this process “rhematization” (Susan Gal and Judith T. Irvine, *Signs of Difference: Language and Ideology in Social Life* [New York, 2019], p. 19). Gal and Irvine present it as an example of what Richard Parmentier has called semiotic “downshifting,” striking a surprising resonance with the conjectural processes performed by pooling, or “subsampling,” layers (Richard Parmentier, *Signs in Society: Studies in Semiotic Anthropology* [Bloomington, Ind., 1994], p. 18.)

The CNN made headlines in 2012 by taking top prize at an annual computer-vision competition.¹⁶⁰ It was not the architecture alone but the unprecedented scale of the winning net—colloquially deemed AlexNet, after its creator—that was noteworthy. Boasting eight layers, five of which were convolutional, and more than sixty million parameters spread across nested pooling layers and the fully connected output, the network’s performance cemented the new paradigm of deep learning. As these indexical algorithms have produced increasingly iconic images, indexical AI has spread, its output ever harder to doubt.

Indexing Language

Nets processing natural language have gained the ability to produce linguistic icons. Operating indexically like all nets, and discovering the heavily indexical aspects of language, these systems have nevertheless demonstrated an ability to produce something like literary language, expression-conveying qualities not found in their parts but only in a holistic fashion. This nascent literary quality might be even more responsible than the achievements of CNNs for our tendency to trust these systems. Recurrent neural nets (RNNs) and their successor, the Transformer Architecture, search for patterns in temporalized sequences. Where the CNN indexicalizes and reassembles an icon, RNNs and

¹⁶⁰ See Alex Krizhevsky, Ilya Sutskever, and Hinton, “ImageNet Classification with Deep Convolutional Neural Networks,” *Communications of the ACM* 60 (May 2017): 84–90.

the Transformer Architecture, which process natural language, produce a potential icon out of a combination of symbols and indexes.

Net-based natural language processing has become “unreasonab[ly] effective,” to modify a phrase that engineers Peter Norvig and Andrej Karpathy separately borrowed.¹⁶¹ Mathematician Eugene Wigner had claimed that mathematics was unreasonably able to explain empirical phenomena, giving us the sense that we “got something out” of the equations that we did not put in.¹⁶² Mathematics, in other words, is a source of discovery, and “unreasonable effectiveness” is a hyperbolic statement of the naïve iconic interpretation of neural nets.¹⁶³

In order to process language (or other sequential data), nets must build an extra set of feedback loops into each calculation layer. Whether a pixel is red or not might be highly dependent on the next pixel over but is usually not dependent on a pixel on the other side of the image. Language, however, displays the opposite tendency. In addition to some

¹⁶¹ See Alon Halevy, Peter Norvig, and Fernando Pereira, “The Unreasonable Effectiveness of Data,” *IEEE Intelligent Systems* 24 (Mar. 2009): 8–12, and Andrej Karpathy, “The Unreasonable Effectiveness of Recurrent Neural Networks,” <http://karpathy.github.io/2015/05/21/rnn-effectiveness/>.

¹⁶² Eugene Wigner, “The Unreasonable Effectiveness of Mathematics in the Natural Sciences,” in *Philosophical Reflections and Syntheses*, vol. 6 of *The Collected Works of Eugene Paul Wigner*, ed. Jagdish Mehra, 8 vols. (Berlin, 1995), p. 540.

¹⁶³ See Rumelhart, Hinton, and Williams, “Learning Representations by Back-Propagating Errors,” p. 536.

short-range dependencies (“Mary runs”) language has long-range dependencies (“Mary, even though she doesn’t want to, runs to maintain fitness”). The conjugation of a verb is a classical indexical component in a sentence, as the third-person singular *s* distinguishes Mary from other possible subjects. These nets have gone one step further and solved some Winograd schema problems, in which the sense of a sentence is dependent on “knowledge of the world.”¹⁶⁴ Terry Winograd’s first example of this sort of problem was the sentence “The city councilmen refused the demonstrators a permit because they advocated revolution.”¹⁶⁵ “They” in this sentence is what Roman Jakobson called a “shifter,” in which meaning “cannot be defined without a reference to the message.”¹⁶⁶ A shifter is an indexical function that conjoins a symbol and an icon (the singular state of affairs “this,” for example). Nets can reduce semantic ambiguity of this kind, linking “they” with the correct referent. But they exploit this capacity to generate much higher-level indexes, ones that connect not only pronouns but also larger units of meaning like genres and styles. Where CNNs indexicalize an icon, nets processing language indexically isolate and generate linguistic icons—meaning in use—as we will see.

¹⁶⁴ Terry Winograd, “Understanding Natural Language,” *Cognitive Psychology* 3, no. 1 (1972): 33.

¹⁶⁵ Ibid. Peirce notes that all pronouns are indexes, since they “syntactically carr[y] the attention *to the word denoting the thing possessed*” (Peirce, “Of Reasoning in General,” in *Elements of Logic*, p. 287).

¹⁶⁶ Roman Jakobson, “Shifters, Verbal Categories, and the Russian Verb,” in *Russian and Slavic Grammar: Studies 1931–1981*, ed. Linda R. Waugh and Morris Halle (New York, 1984), p. 42.

All nets processing language use some form of embedding to make strings of words into data input. Algorithms like word2vec embed words by assigning them to tokens (vectors, in this case) to input into a net. Initial vectors are random, and the algorithm “learns” that “words that appear in similar contexts have similar meanings” (*D*, p. 181). These vectorized words then become the input for neural nets, which concatenate these meanings into syntactically correct, semantically approximate expressions. Tomas Mikolov and his colleagues express surprise at the ability of their word-vectors to learn analogy: “The result of a vector calculation $\text{vec}(\text{‘Madrid’}) - \text{vec}(\text{‘Spain’}) + \text{vec}(\text{‘France’})$ is closer to $\text{vec}(\text{‘Paris’})$ than to any other word vector.”¹⁶⁷ These embedding algorithms have managed to produce the prediction “queen” in answer to the query: “king – man + woman.” The autonomy of such a semantic space resembles what Jakobson called “contiguity disorder,” a type of aphasia characterized by “agrammatism.”¹⁶⁸ For Jakobson, this disorder means that the patient possesses metalanguage but no object language—they can substitute meanings for one another but cannot synthesize sentences that rely on context. They understand “gestalt” but are trapped in its paradigm.¹⁶⁹ As in word2vec, the metaphoric pole of

¹⁶⁷ Tomas Mikolov et al., “Distributed Representations of Words and Phrases and Their Compositionality,” *Advances in Neural Information Processing Systems* 26 (2013), p. 1, <http://papers.nips.cc/paper/2013/file/9aa42b31882ec039965f3c4923ce901b-Paper.pdf>.

¹⁶⁸ Jakobson, “Two Aspects of Language and Two Types of Aphasic Disturbances,” in Jakobson and Morris Halle, *Fundamentals of Language* (New York, 1971), pp. 85, 86.

¹⁶⁹ *Ibid.*, p. 87.

language becomes a series of proximities without syntax. Word embedding produces this semantic space, and the net is left to learn the rules of contiguity.

The RNN functions on a simple premise: at each layer, the output must contain both the normal layer forwarding (activation, weights, input values, and bias) and a hidden state that carries information forward from earlier in the layers—context, in the form of a vector storing previous information in the sequence. The recurrence is a *for loop*, in programming terms, iterating through the string until all relevant dependencies are learned. But the memory buffer will slowly minimize lower probabilities out of existence, a problem called the vanishing gradient. The addition of “long short-term memory” was meant to allow RNNs to index longer-range dependencies.¹⁷⁰ This architecture differentiates the memory buffer by making it into a cell composed of a forget gate that pushes some values near 0 and not does not pass them forward (a mathematical index).¹⁷¹ The LSTM generates a kind of indexical heatmap of a text, a linguistic version of the feature maps produced by convolutional layers. But the LSTM can only reduce, not eliminate, the vanishing gradient of the RNN, leading to their replacement, in 2017, by the “Transformer architecture,” which relies on “self-attention” to relate “different positions of a single sequence in order

¹⁷⁰ See Sepp Hochreiter and Jürgen Schmidhuber, “Long Short-Term Memory,” *Neural Computation* 9 (Nov. 1997): 1735–80.

¹⁷¹ Chollet points out that the names of the gates are over-interpretation because the weights actually determine the functions during the training; see Chollet, *Deep Learning with Python*, p. 204.

to compute a representation of the sequence.”¹⁷² Rather than storing information through the forward pass, attention selects for relevant information at each node, imitating the indexical function of consciousness.

Paying attention is patently indexical. Think of the ubiquitous banner messages that force our eyes to our screens; these are indexes driving the attention economy.¹⁷³ When we pay attention, we choose an anchor for our representations and allow the mind to range over qualities, relations, and other signs around that center. Peirce notes that the “subject of a proposition” is an attentional index, in this sense, communicating a “quality” against which the predicate can be checked.¹⁷⁴ “Self-attention” performs just this operation.

The Transformer contains three separate attention mechanisms each with multiple attentions functions (multihead): one for its encoder, one for its decoder, and one for their interaction called the “encoder-decoder” (Figure 15, 16).¹⁷⁵

¹⁷² Ashish Vaswani et al., “Attention Is All You Need,” *Advances in Neural Information Processing Systems* 30 (2017), pp. 9, 2, <http://papers.nips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.

¹⁷³ “Anything which focusses the attention is an index” (Peirce, “Genuine and Degenerate Indices,” in *Elements of Logic*, p. 285).

¹⁷⁴ Peirce, “Supplement of 1893,” in *Elements of Logic*, p. 428.

¹⁷⁵ Vaswani et al., “Attention Is All You Need,” pp. 3–5.

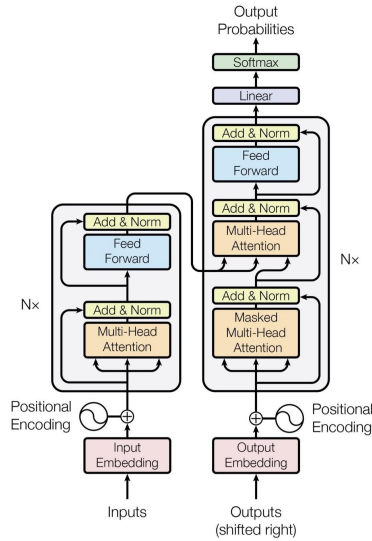


Figure 1: The Transformer - model architecture.

Figure 15. The Transformer architecture.




Figure 16. Multihead attention.

Let's say you have a query: What should the next word be in a sentence? What is the correct word at a given position for a translation of a sentence from French to English? Attention takes two input vectors—values and keys, where keys are indexed to the values,

such that they are often initially equivalent—and multiplies the query vector by the specific key, which produces a set of scores for the k vector, establishing “how much each word will be expressed at this position” (Figure 17).¹⁷⁶ For technical reasons, all these products will fall between 0 and 1. The attention head then uses an exponentialization function called the “softmax,” which creates two groups of the product values, one very close to 0 and one very close to 1 (Figure 18).¹⁷⁷ The softmax mathematically indexes a set of relevant units. These products are then multiplied by the values vectors and summed to produce the result of this attention. Once all three attention heads have acted, a prediction is created on the basis of a composite sense of relevance. The multihead approach, which launches the attention calculation into very high dimensions, has allowed the Transformer to learn both semantic space and syntactic order.

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

Figure 17. Softmax in the attention equation.

¹⁷⁶ Jay Alammar, “The Illustrated Transformer,” <http://jalammar.github.io/illustrated-transformer>.

¹⁷⁷ Ibid.

Scaled Dot-Product Attention




Figure 18. *Softmax in the attention mechanism.*

The old suspicion that digital computers are syntax machines with no capacity for meaning should finally be laid to rest. These machines do not suffer from Jakobson’s other aphasic pathology, “similarity disorder,” never knowing content but only order. Jakobson writes that patients lacking metaphor and possessing only metonymy “fail to shift, as Peirce would say, from an *index* or *icon* to a corresponding verbal *symbol*.”¹⁷⁸ In other words, they fail to attend to the shifts in symbol-icon connections, a form of semantic creep.¹⁷⁹ The Attention mechanism uses the softmax index to find patterns among the

¹⁷⁸ Jakobson, “Two Aspects of Language and Two Types of Aphasic Disturbances,” p. 80.

¹⁷⁹ On iconicity in language, see Winfried Nöth, “Semiotic Foundations of Iconicity in Language and Literature,” in *The Motivated Sign*, vol. 2 of *Iconicity in Language and Literature*, ed. Olga Fischer and Max Nänny, 18 vols. (Philadelphia, 2001), pp. 17–28.

symbols that correspond to redundancies in usage that we find meaningful or construe as icons.

Language models based on the Transformer architecture have been able to generate philosophical essays about themselves, consistent if imperfect poetry in various author styles, and even to write code (or what looks like code) on command.¹⁸⁰ These large-scale models create composite indexes based on high-dimensional multihead attention mechanisms that trace the distance between very complex qualities in language, allowing them to communicate firsts or present us with potential icons.¹⁸¹ At scale, these potential icons seem to carry their own rhetorical force, pushing adoption of the models in spite of their risks.¹⁸²

¹⁸⁰ See Tom B. Brown et al., “Language Models Are Few-Shot Learners,” *ArXiv.org*, <http://arxiv.org/abs/2005.14165>.

¹⁸¹ The models are based on “generative pre-training” (GPT) on a large amount of unlabeled text, as much as nearly a trillion words producing 175 billion parameters in the net (ibid., p. 73). The large amount of data contains “naturally occurring demonstrations of many tasks across diverse domains,” allowing these systems to respond to natural-language queries and learn very quickly on the basis of a few examples (“few-shot”) Alec Radford et al., “Improving Language Understanding by Generative Pre-Training,” 2018, www.cs.ubc.ca/~amuham01/LING530/papers/radford2018improving.pdf.

¹⁸² See Gebru et al., “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” p. 14.

Take the celebrated article one such model produced (Figure 19).¹⁸³ On a premise similar to that of Arthur Conan Doyle’s *The Lost World* (1912), the algorithm was able to generate a formally correct news article, respecting the generic requirements with only minor flaws (dialect or dialectic). Successes like this have led to debates about these systems has focused on whether they have become intelligent by learning contextual or worldly information from these very large language datasets, a notion that falls squarely into the naïve iconic interpretation of AI, failing to connect architecture and expression in concrete analysis.¹⁸⁴

¹⁸³ See “Better Language Models and Their Implications,” *OpenAI*, 14 Feb. 2019, <http://openai.com/blog/better-language-models/>.

¹⁸⁴ Yannic Kilcher’s account is among the most sober and technically interesting; see Yannic Kilcher, “GPT-3: Language Models Are Few-Shot Learners (Paper Explained),” YouTube, 29 May 2020, www.youtube.com/watch?v=SY5PvZrJhLE&ab_channel=YannicKilcher. See also Justin Weinberg, “Philosophers On GPT-3 (Updated with Replies by GPT-3),” *Daily Nous*, 30 July 2020, <http://dailynous.com/2020/07/30/philosophers-gpt-3/>. The systems’ relatively subpar performance on NLI (natural language inference) tasks, which most directly involve problems that can’t be solved by reference to language alone, suggests the limits on even this very large scale of language data.

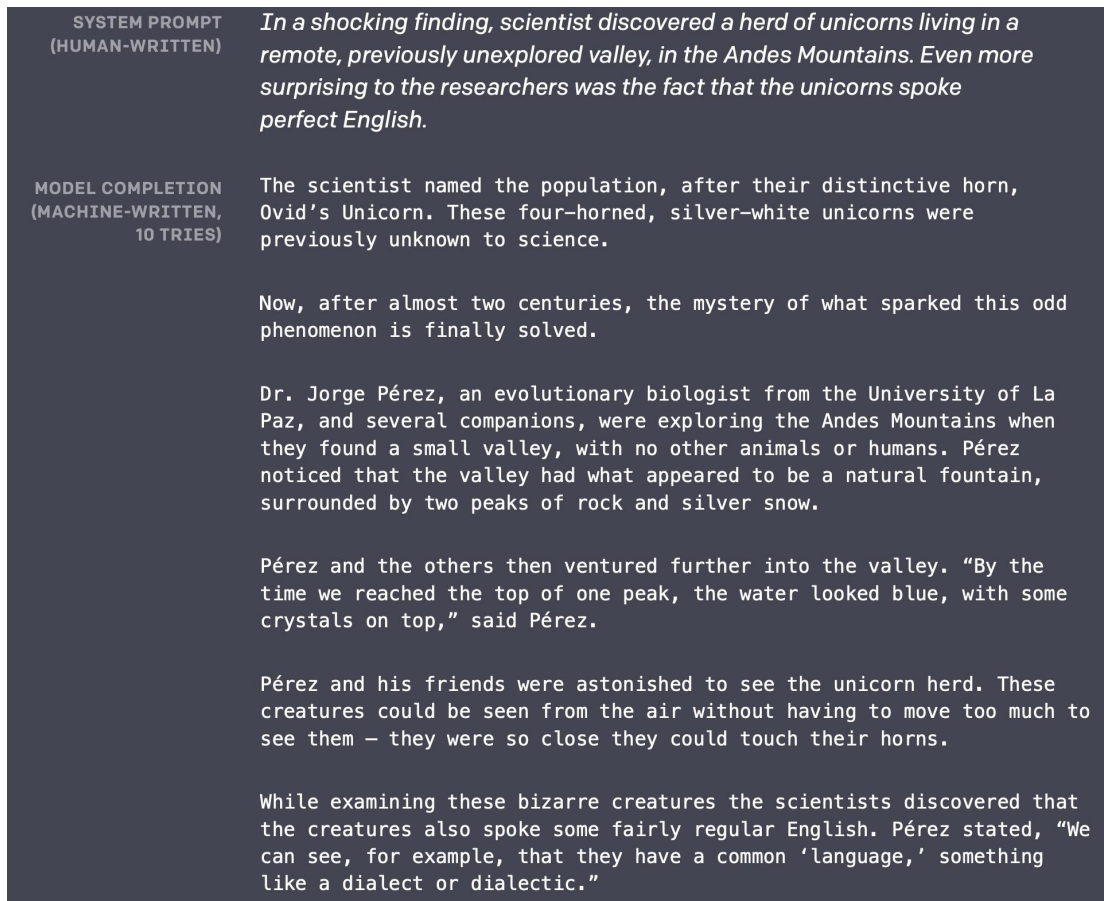


Figure 19. An article generated by GPT-2 in the style of science journalism.

When the article attributes four horns to the eponymously single-horned creatures, the error is an index that dislocates us from the universe of discourse in which the article otherwise immerses us. Whether science fictional, fantastical, or as narrow as binary logic, universe of discourse require “an index in the environment common to speaker and auditor.”¹⁸⁵ This index—computer science calls it “\interoperability, and structuralism

¹⁸⁵ Peirce, “Universe of Discourse,” in *Elements of Logic*, p. 536.

calls it the phatic function—must exist because universes “cannot be described.”¹⁸⁶ To speak of unicorns is to locate ourselves in a universe we fall out of if they are ascribed four horns. Semantic entities like unicorns are not grounded only in the mental picture of a flying horned horse. They possess, as part of their meaning, an indication of potential universes to which they belong, a kind of semantic-indexical embedding. The unicorns article barely fails to immerse us in such a universe (indeed, we have never seen the contradiction of the horns pointed out in the breathless literature that followed on this article).

Large-scale language models arrive at the synthesis of text indexically, as they rely on the attention mechanism. But they can communicate a first, a quality—in this case, genre. What impresses the reader of the article is that it is a news article, that it appears as science reporting. In other words, the attention mechanism drawn out over a very large corpus is able to produce generically valid expressions, an instance of something we immediately recognize but which was produced by billions of indexical parameters operating on trillions of symbol instances. The naïve iconic interpretation comes from this feeling of recognition, taking the quality as the sign. Large-scale language models generate powerful linguistic icons from high-dimensional digital indexes. These icons are not just

¹⁸⁶ Peirce, “Universes and Predicaments,” in *The Simplest Mathematics*, vol. 4 of *The Collected Papers of Charles Sanders Peirce*, p. 544. See also Jakobson, “Two Aspects of Language and Two Types of Aphasic Disturbances,” p. 90.

orderings of words that resemble the flow of thought.¹⁸⁷ They are closer to W. J. T. Mitchell's account of the “verbal icon,” communicating singular linguistic qualities, not limited to verbal pictures, but also including genre and mood.¹⁸⁸ To the extent that neural nets can generate such qualities, they participate in the literary aspect—what Jakobson called the “poetic function”—of language, where language's object domain is language itself.¹⁸⁹ This reflexive appearance underwrites both the adoption of these systems and the naïve iconic interpretation of their abilities. We are in the process, not yet complete, of plugging these powerful index machines into our social-conceptual icons, allowing their expressions to adopt the ring of truth, the benefit of the doubt. Indexical AI lacks icons of its own, but it can generate signs we are prone to take as sources of discovery.

Indexical Labor and Algorithmic Capital

Indexical AI has become a “general condition of production,” linking the moving parts of the economy to each other.¹⁹⁰ The platform economy is not limited to targeted ads on

¹⁸⁷ See Peirce, “Universes and Predicaments,” p. 544.

¹⁸⁸ W. J. T. Mitchell, *Iconology: Image, Text, Ideology* (Chicago, 1986), p. 25. The term *verbal icon* is W. K. Wimsatt and Beardsley's, who note that this trope is “not merely a bright picture (in the usual modern meaning of the term image) but also an interpretation of reality in its metaphoric and symbolic dimensions,” underscoring the ability of linguistic nets to convince (William K. Wimsatt and Monroe C. Beardsley, *The Verbal Icon: Studies in the Meaning of Poetry* [Lexington, Ky., 1954], p. x).

¹⁸⁹ Jakobson, “Linguistics and Poetics,” *Language in Literature*, ed. Krystyna Pomorska and Stephen Rudy (Cambridge, Mass., 1987), p. 69.

¹⁹⁰ Nick Dyer-Witheford, Atle Mikkola Kjosen, and James Steinhoff, *Inhuman Power: Artificial Intelligence and the Future of Capitalism* (London, 2019), p. 52.

social media, as important as these are for the indexing of consumer desires and the alteration of production. We drive using nets, taking unfamiliar routes dictated by Google data. The prices we see on Amazon, which increasingly pressure offline prices too, are probably the result of personalized net-based fluctuations.¹⁹¹ When you ask a rideshare to take you to the airport, the matching algorithm points driver and rider to each other and points them through space to a semantified goal (such as “terminal C”). When we follow Google’s routes, follow through on an Amazon recommendation, or confirm a time and location for car pickup, we enrich these companies’ indices. We even add data capital to these companies when we deviate from their routes, decline their recommendations, or cancel a pickup. More data means more flexible and powerful algorithmic indexes. This semiotic infrastructure bleeds upward into the daily interpretive tasks we face in the digital world. The digital index manipulates and shadows identity,¹⁹² while subsuming labor in the platform economy.¹⁹³ AI links our daily hermeneutic activity to global logistics by means of the index, tightening the loop of desire and understanding with production and distribution.¹⁹⁴ The naïve iconic interpretation of the net is virtually forced on us—resistance to hype tends to be futile.

¹⁹¹ Lina M. Khan, “Amazon’s Antitrust Paradox,” *The Yale Law Journal* 126 (Jan. 2017), www.yalelawjournal.org/note/amazons-antitrust-paradox

¹⁹² See Antoinette Rouvroy, Thomas Berns, and Elizabeth Libbrecht, “Algorithmic Governmentality and Prospects of Emancipation,” *Revue* 177 (Oct. 2013): 163–96.

¹⁹³ See Nick Srnicek, *Platform Capitalism* (Malden, Mass., 2017).

¹⁹⁴ See Brian Justie, “Little History of CAPTCHA,” *Internet Histories* 5 no. 1 (2020): 30–47.

Underlying this semiotics, and the daily hermeneutics it forces on us, is a reordering of the elements of capital. What appears to be a technology of distribution now shapes both labor conditions and capital itself. The index that appears to us as a point, a context, or a location is the channel along which labor flows. The partly automated data flow between those channels and the allocation of the resources of enterprise make up the structure of contemporary capital, an attractive force organizing the chaos of digital signs. We formulate this conclusion as two theses for future work: (1) indexical AI valorizes data by subsuming labor into the algorithmic index of capital; (2) the shape of capital is a feedback loop, partially automated by indexical AI, among context and action, data and judgment.

Hito Steyerl has identified the condition of ubiquitous ambient digital semiotics as “free fall,” in which grounding is only possible as if from above, so that data panoramas “do not actually portray a stable ground.”¹⁹⁵ But even if free fall continues indefinitely, it must be *falling* for the metaphor to work. The gravity in digital free fall is the index specific to contemporary AI. As we enter the age of ever bigger data, learning systems have yawned into existence as mega indexes tracing the outline of capital today.

¹⁹⁵ Hito Steyerl, “In Free Fall: A Thought Experiment on Vertical Perspective,” *e-flux* 24 (Apr. 2011), www.e-flux.com/journal/24/67860/in-free-fall-a-thought-experiment-on-vertical-perspective/.

3. “The Nonmachinables”¹⁹⁶

In the business of logistics, the “last mile” problem refers to the final stages of transporting an item from point-a to point-b, which often tends to be the most expensive, laborious, and occasionally dangerous.¹⁹⁷ Negotiating the last mile in logistics, like overcoming the last percentile of error in a machine-learning model, as discussed in the preceding chapter, is a messy, high stakes business. However, in both cases, these are not simply technical hurdles to be overcome by sufficiently advanced technologies, but rather are ongoing matters of political-economic negotiation. It is here, within the figurative space of the “last mile,” where the concept of asymptotic labor is rendered most tangible, as will be shown in the following study of the “hidden” workforce that toils away behind the scenes at United States Postal Service.

¹⁹⁶ Justie, Brian. “The Nonmachinables.” *Logic Magazine*, 2021. <https://logicmag.io/distribution/the-nonmachinables/>.

¹⁹⁷ Alimahomed-Wilson, Jake. “The Amazonification of Logistics: E-Commerce, Labor, and Exploitation in the Last Mile.” In *The Cost of Free Shipping*, edited by Jake Alimahomed-Wilson and Ellen Reese, 69–84. Amazon in the Global Economy. Pluto Press, 2020. <https://doi.org/10.2307/j.ctv16zjhj.11>. Altenried, Moritz. “On the Last Mile: Logistical Urbanism and the Transformation of Labour.” *Work Organisation, Labour & Globalisation* 13 (April 1, 2019): 114–29. <https://doi.org/10.13169/workorglaboglob.13.1.0114>.

This chapter began as a short ethnographic undertaking, based on my time spent at the Remote Encoding Center in Salt Lake City, UT in the summer of 2020, where this specialized workforce is stationed. I was able to interview a number of workers and managers about their work as well as the history of the REC, and spent several hours observing the work of keying itself. However, it became abundantly clear that the peculiarities of this line of work could not be conveyed successfully without proper historicization, chronicling a century-and-a-half of technological development and strife between labor and management that necessitated the work performed by keyers. Accordingly, I conducted considerable historical research, primarily in postal archives and congressional records, in order to better piece together this tumultuous trajectory, and better situate my observations at the REC. From simple, mechanical apparatuses to help expedite the canceling of postage stamps in the 19th century, to the flashy public-private partnership that exists today between the USPS and NVIDIA, the leading producer of computer chips powering the ongoing machine learning boom, new technologies have been unilaterally rolled out by postal management and their congressional overseers, as a means of indirect labor relations.¹⁹⁸

¹⁹⁸ Merritt, Rick. “Sharpening Its Edge: U.S. Postal Service Opens AI Apps on Edge Network.” NVIDIA Blog, May 6, 2021. <https://blogs.nvidia.com/blog/2021/05/06/edge-ai-usps/>.

The Bureau of Hards

Sometime around 1870, the New York City Post Office established a new department, staffed by a small team of specialists. According to an 1871 profile in *Harper's Magazine*, these postal workers spent their days scrutinizing what looked like “the footprints of a gigantic spider that had, after wading knee-deep in ink, retreated hastily” across envelopes and postcards.¹⁹⁹ In reality, these would-be arachnologists were employed to make sense of the “miserable chirography” of city residents, whose poor penmanship was causing unacceptable delays in delivery.

One expert decipherer recalled working on a letter that had arrived back in New York after traveling hundreds of miles over four days, repeatedly rejected as illegible by clerks in regional post offices. He studied the chickenscratch for a full workday before finally cracking the code: it was addressed to Chappaqua, a city just thirty miles north. Taking its name from the informal term clerks reserved for these most challenging pieces of mail, the new outfit came to be known as the Bureau of Hards.

Delivering hards, no matter the cost, is a reflection of the US Post Office's commitment to truly universal service—a radical vision of democratic communications infrastructure

¹⁹⁹ Harper's New Monthly Magazine. “New York City Post-Office,” November 1871. 662.

enshrined in the Post Office Act of 1792.²⁰⁰ No matter the sender, the recipient, or the distance separating origin and destination, federal code stipulated that the Post Office must “bind the nation together.” As Alexis de Tocqueville put it in his 1835 treatise *Democracy in America*, the US mail system, unlike its European counterpart, “was organized so as to bring the same information to the door of the poor man’s cottage and to the gate of the palace.”²⁰¹ To live up to this idealistic ethos, hards must be treated no differently than easies.

But, as the errant letter destined for Chappaqua demonstrates, universal mail service tends to be extremely laborious. Supplanting human postal workers—slow, error-prone, and wage-requiring—with nonhuman proxies has long been a prospect with considerable purchase for postal management. The first machines arrived in post offices in the 1870s, and it’s no coincidence that the first postal worker unions were formed then, too. By the turn of the twentieth century, the Post Office and its governmental overseers had set into motion an unceasing drive to maximize the role of machines and minimize the role of humans.

²⁰⁰ Post Office Act of 1792,

<https://collections.si.edu/search/record/ark:/65665/hm8c0ae955a6e4f445e93c830f692e2c1b4>. See also: Fuller, Wayne. *The American Mail: Enlarger of the Common Life*. Chicago: University of Chicago Press, 1972.

²⁰¹ Tocqueville, Alexis de. *Democracy in America*. Translated by Harvey C. Mansfield and Delba Winthrop. 1st edition. Chicago, Ill.: University of Chicago Press, 2002.

Today, the United States Postal Service possesses the third-largest information technology infrastructure in the world—a rarely cited superlative.²⁰² Not often included in the discourse around Big Tech, the USPS controls a sophisticated and sprawling computer network, linking together over 30,000 facilities and nearly 10,000 pieces of automated machinery, shuttling 150 billion pieces of mail per year between 150 million delivery points: houses, businesses, and PO boxes from Utqiagvik, Alaska, to Key West, Florida. The number of mail pieces per postal worker, a rough measure of automation’s impact, has more than doubled since 1950.

Despite significant advances in postal technology made since the mid-twentieth century, however, the USPS remains the country’s largest public-sector employer. The majority of its workforce is stationed at the input or output stage of what is, in effect, an enormous circuit: clerks are responsible for getting mail into the mailstream, and letter-carriers handle mail once it has exited. No single mandate better captures the thrust of modern postal operations than that of realizing a fully automated mailstream capable of connecting clerk and carrier with zero intervention from humans along the way.

²⁰² United States Postal Service. “Information Security.” Information Technology. Accessed March 2, 2023. https://about.usps.com/strategic-planning/cs10/CSPO_12_2010_FINAL_054.htm,

But a tiny fraction of the USPS's half a million workers—about two-tenths of 1 percent— toil away in a modern day Bureau of Hards. Two miles south of the Salt Lake City International Airport, in a drab warehouse, these workers parse the squiggled and smudged addresses emblazoned on each piece of mail that has proven illegible to the advanced machine-readers deployed in processing plants across the country. The 1,100 workers on staff at the Remote Encoding Center (REC) tend to the “nonmachinable” scraps discarded by the Postal Service's automated leviathan. They ensure that more than one billion pieces of mishandled, misdirected, and misidentified mail arrive at their destinations each year.

Hards have never been simply a technical problem in need of a technical solution. Rather, hardness is better understood as an index of the social and political conditions under which mail is delivered. Taken together, these two deciphering operations—the nineteenth-century Bureau of Hards and the twenty-first-century REC—become legible as something like the origin and destination of an arduous and ongoing struggle between postal management and postal workers over the question of technological change.

RIPS 122K

Workers stream into and out of the Salt Lake City REC at all hours, with shifts staggered to begin every 15 minutes, 24 hours per day, 365 days per year. Once inside, Data Conversion Operators, informally known as keyers, consult one of the many large monitors

scattered around the facility displaying a clutter of acronyms and numbers. On the morning when I visited the REC last summer, this coded to-do list read “PARS 209K / RIPS 122K / PRES 28K,” indicating that just over 350,000 items were awaiting judgment from a discerning human eye (Figure 20.)

At the center of the 77,000-square-foot facility is an elevated platform that keyers cheekily refer to as “air-traffic control.” From this vantage, managers consult a comically large array of computer screens with charts, spreadsheets, and data visualizations that track the incoming and outgoing flow of mail. What arrives at the REC is not the physical piece of mail itself, but rather a digital surrogate. Each time a sorting machine in a postal plant encounters an address it cannot match with one in the USPS’s database, it takes a snapshot and automatically sends it on a virtual detour to the REC. Keyers receive these images in large batches and make quick work of them; eyes quickly scanning for visual clues, fingers dexterously entering relevant address data.

Once the address data has been entered into the keying interface, the image is sent back to the plant, where the sorting machine applies a new barcode to the parcel, ensuring it remains legible to the subsequent sorting machinery it will encounter en route to its destination. This seemingly simple image-processing task depends upon a dense and dated patchwork of software and hardware linking the REC to more than 300 postal hubs across the country.

Acronyms abound at the REC, indicating the sheer number of image-processing systems in use. Different sorting machines, designed to handle different types of mail, send different types of data to the REC, using different software platforms designed over the last several decades. Some links in this complex chain have been designed in-house, but many have been farmed out to contractors, including Lockheed Martin and Siemens.

The goal of consolidating these many overlapping but incompatible systems has long been a high priority for keyers and management alike. RIPS, which launched in 2019, is the newest consolidation effort. Once completed, it will funnel the data from IPS, PICS, FICS, and PRES—four currently siloed systems that handle letters, packages, “flats” like magazines, and other postal paperwork—into a single keying interface.

In the past, different sections of the facility were hardwired to handle these different types of mail, creating a bustling, sometimes chaotic environment. Keyers would walk, or occasionally run, from one section to another as keying queues ebbed between parcel types throughout their shift. Endcaps on cubicles still display “*No more cutting through aisles!*” signs, despite the REC’s current library-like stillness. Much to management’s delight, less time spent on foot means more time spent keying.



Figure 20. Remote Encoding Center, Salt Lake City, UT.

Every keying station is stocked with a keyboard, a monitor, and three or four desktop towers, each dedicated to one of the many acronymic systems in use. Many of these towers are brand new, but run as virtual machines emulating legacy software platforms developed decades ago. In lieu of physically moving throughout the facility, keyers speed back and forth between different programs every several minutes—say, from APPS (packages) to PARS (change-of-address forms) and back to APPS—each time reentering their username and password.

The work of keying is almost unfathomably fast-paced. But it is graceful, not frenetic. Most keyers sport headphones, listening to music or podcasts as they nimbly flit through an unending sequence of rasterized black-and-white images, logging upwards of 10,000 keystrokes per hour. Some have elected to work at standing desks, a benefit associated with an intensive ergonomics program won by their union. To speed up the information exchange between processing plants and the REC, the images are heavily compressed, producing low-resolution depictions of physical objects that are, oftentimes, in far from pristine condition. Torn shipping labels, waterlogged envelopes, and smeared ink are common. This combination makes for an oddly compelling aesthetic, somewhere between the warped scanlines of artist Bruno Munari's *Xerografia* and the lo-fi letterforms of mid-aughts reCAPTCHA puzzles.

What scant media coverage the REC has received has almost universally focused on the idiosyncrasies of bad handwriting. The thousands of letters addressed to SANTA, NORTH POLE sent each year by grade-schoolers still honing their penmanship is a recurring motif.²⁰³ One keyer reminisced about envelopes decorated with hand-drawn pin-ups, barbed wire, and skulls that inmates at a local prison used to send, and which he would occasionally receive for keying in the 1990s.

²⁰³ Opam, Kwame. "Postal Service Takes 'Operation Santa' Letter Campaign National." *The New York Times*, November 16, 2020, sec. U.S. <https://www.nytimes.com/2020/11/16/us/usps-operation-santa.html>.

But according to several keyers I spoke with, the lion’s share of the five million items that pass through the REC each day are not handwritten addresses. Contemporary machine-readers, it turns out, can read handwriting with relative ease, leaving keyers to trudge through a bottomless pile of machine-printed detritus, much of it cheaply printed junk mail slung by mass marketeers.

Crash Program

The vast infrastructure required to affix problematic parcels with a packet of human-verified metadata—a thankless clean-up job, performed in the service of machines—is the culmination of a century and a half of technological change. This arc has not been one of linear progress, but rather one of tumult and negotiation, as postal workers from the late nineteenth century onward fought to retain autonomy in the face of encroaching machinery.

The first to arrive was the mechanical “canceller,” a device patented in 1876 by a pair of Boston inventors.²⁰⁴ The Post Office Department—not yet the United States Postal Service—contracted the production of one hundred cancellers, and allocated them to the nation’s busiest post offices to assist clerks in the slow-going process of manually voiding

²⁰⁴ Scheele, Carl H. *A Short History of the Mail Service*. Smithsonian Institute Press, 1970.

postage stamps to prevent repeated usage. The contraption's hand crank rapidly fed letters through a pair of rollers, allowing clerks to cancel fifteen times more postage per hour than was possible by hand. Mechanical cancellers, like many of the technological novelties that would eventually make their way into the post offices, helped to set unprecedented expectations for postal worker productivity.

Innovations like this helped the Post Office keep up with a dizzying uptick in mail volume. Sending and receiving mail had become gradually more accessible during the nineteenth century, as a growing share of the public was now within spitting distance of a local post office, and a new policy offered rural delivery for no additional charge. A new class of senders had also entered the scene. Just a few years before the mechanical canceller was introduced, Montgomery Ward had sent its first mail-order catalog, using the Post's unparalleled delivery network. Mail volume ballooned by a factor of fifty in the decades leading up to the twentieth century, and showed no signs of slowing down.

The rising demands of commercial mailers put even greater strain on postal operations in the first half of the twentieth century. The cumbersome work of sorting the mail was still performed using the peek-and-poke method: clerks would glance at and then manually deposit each item, one by one, into an appropriate cubbyhole. This outmoded process, developed in the eighteenth century, was plainly incapable of scaling up to the degree necessary. By midcentury, a full nickel's worth of each six-cent stamp still went toward

covering the labor costs of sorting the mail. With competition mounting from private upstarts like the quickly expanding United Parcel Service, the Postmaster General publicly committed the department to a “crash program of modernization and mechanization.”²⁰⁵

In 1957, the Transorma, an impressive piece of Dutch engineering that mechanized the peek-and-poke, was brought stateside. Five clerks at a time would take their posts at terminals within the belly of this hulking, fifteen-ton apparatus. Letters were shuttled through the machine’s guts, briefly pausing in front of a clerk who would manually punch in a memorized code pegged to the letter’s destination. The machine would then whisk the letter away to a bin for subsequent processing. This new breed of machine enabled a five-fold increase in sorting productivity, but also gave management greater control over workflows.

By the 1960s, utility bills, catalogs, advertisements, invoices, receipts, and other forms of impersonal, bulk communication had come to account for more than 80 percent of all mail—clogging up the mailstream, but also providing a critical revenue stream. It was clear to postal management that, on their own, hardware innovations like the Transorma

²⁰⁵ Rubio, Philip F. *Undelivered: From the Great Postal Strike of 1970 to the Manufactured Crisis of the U.S. Postal Service*. UNC Press Books, 2020.

would not be enough. Accordingly, in 1963 the Post Office introduced what might be best understood as its first innovation in software: the five-digit ZIP code.

Above all, ZIP codes served as a new standardization protocol, transforming an unruly map into an efficient mosaic. Encoded in each five-digit string was a surfeit of data, helping to direct each parcel through a carefully delineated geographical hierarchy, from regional processing plant down to localized delivery zone. Not only did this numerical logic significantly simplify manual mail sorting, it also greased the skids of mechanization.

Homegrown alternatives to the Transorma were developed throughout the 1960s, designed specifically to take advantage of the new ZIP system, which theoretically enabled faster keying by workers. As installation expanded, sorting machines began to play the part of crucible for a brewing hostility between postal workers and management. By 1968, the Post had purchased 145 Multiple Position Letter Sorting Machines (MPLSM) designed by the Burroughs Corporation, famous for its hand-crank calculators. Unlike the Transorma, which advanced to the next letter only after the clerk had entered a code, the new American-built MPLSM made pacing a point of contention. Workers wanted to be able to advance the sorting machine themselves, letter by letter, allowing for more flexibility and precision in the keying process. Management wanted to program the machine to advance at automatic intervals, maximizing productivity and ensuring predictable throughput.

The distinction between operator-pacing and machine-pacing was the subject of considerable research: Which was more efficient? More sustainable? More cost effective? Consultants were hired to conduct extensive psychophysical studies, monitoring eye movements and keystrokes, fatigue and focus, hoping to determine the optimal balance between speed and accuracy. Despite the findings in these reports, management opted for machine-pacing, disregarding the ample evidence that this would greatly reduce overall efficiency and further degrade working conditions.

The tug-of-war over the MPLSM was not an isolated incident. Grievances over wages and working conditions—facilities were dated and deteriorating, the hours were getting longer, the productivity quotas higher—were piling up in postal facilities across the country. But disputes over the role of technology in particular helped to set the stage for, and ultimately played a starring role in, the most significant reshaping of the Postal Service in its history.²⁰⁶

²⁰⁶ Ellis, Ryan. *Letters, Power Lines, and Other Dangerous Things: The Politics of Infrastructure Security*. MIT Press, 2020.

Processing Progress

In the spring of 1970, several thousand disgruntled postal workers in New York City walked off the job.²⁰⁷ Over the next week, they were joined by over 200,000 of their colleagues around the country, forming the largest wildcat strike in American history, and bringing nationwide postal operations to a near standstill. This action put immense pressure on management and congress to come to the bargaining table. The hardscrabble negotiations that ensued between labor, management, and policymakers carved a new route for postal operations—a route leading directly to the REC.

These negotiations resulted in the 1970 Postal Reorganization Act (PRA), signed by President Richard Nixon, which minted the United States Postal Service. It earned postal unions the right to collectively bargain over wages, benefits, and working conditions for the first time, something expressly prohibited for the Post Office Department, which had been a part of the US Cabinet. This apparent win for organized labor, however, came at a cost, as the PRA also cemented a new ideological foundation undergirding all postal operations.²⁰⁸ It made manifest the decree of a federal commission assembled in 1967 and

²⁰⁷ Rubio, Philip F. “After the Storm: Postal Politics and Labor Relations Following the 1970 U.S. Postal Wildcat Strike, 1970–1981.” *Employee Responsibilities and Rights Journal* 30, no. 1 (March 1, 2018): 65–80. <https://doi.org/10.1007/s10672-017-9303-7>.

²⁰⁸ Baxter, Vern K. *Labor and Politics in the U.S. Postal Service*. Plenum Studies in Work and Industry. Springer, 2013.

helmed by an ex-chairman of AT&T, one of the USPS's key competitors in the private sector: "Today the Post Office is a business."

The PRA renewed the Post's commitment to provide "a basic and fundamental service," but made clear that a balanced budget was of equal—perhaps, greater—importance. But in the eleventh hour of negotiation, an important caveat to this prioritizing of fiscal concerns was hashed out. Congressional representatives had attempted to slip in one last amendment stating that the "Postal Service shall promote modern and efficient operations and should refrain from [any activity] which restricts the use of new equipment or devices." If accepted, this would significantly erode labor's bargaining power by letting postal management make unilateral decisions about technological changes. But labor representatives refused to concede on this point. Tech, they maintained, *must* be bargainable.²⁰⁹

After the amendment was struck down, the battles over technology continued with a renewed vigor, centered around the automated equipment that had begun to replace the mechanized machinery of the 1950s and '60s. The key difference between the two paradigms lay in the question of who—or *what*—would be responsible for the work of actually reading the address line, a necessary first step before any sorting could commence.

²⁰⁹ Ibid.

The promise of postal automation, which would require delegating the reading to machines, had long been undermined by optical-character recognition (OCR) technology's failure to deliver on its own promise. OCR developers had a track record of lofty assurances about the efficacy of their machine-reading systems, stretching back to the early patents filed by AT&T in the 1920s. While the fanfare around OCR was clearly overblown, it had created a deluge of commercial research and investment into the technology in the 1950s.

The Post Office Department had begun experimenting with OCR in the late 1960s, but had quickly run aground. Throughout the 1970s, various attempts were made to enhance existing Multiple Position Letter Sorting Machines by integrating an OCR that could identify the bottom-most line of an address and read the five-digit ZIP code. All mail fed into this OCR first had to be examined and presorted by workers on site, as these initial iterations could only read a fraction of the most popular typefaces in use. They were also highly susceptible to paper jams, and handwritten mail remained especially elusive. Consultants determined that until 85 percent of mail could be accurately read by OCR, the mechanized MPLSM—and its waged, unionized operator—would be both more efficient and more cost effective than the automated alternative.

Automation would become viable in Reagan's 1980s, which were a hotbed of innovation for both postal technology and management-labor relations. The decade opened with an

unprecedented showing of strikebreaking force, as Reagan fired over 10,000 air-traffic controllers who were protesting over wages and working conditions. This sent a message to the American Postal Workers Union (APWU) who, weeks prior, had been on the verge of calling for another large-scale strike. Reagan's airport intervention spooked the postal unions, who canceled their pending strike authorization vote. The militant worker energy that came to define the 1960s and '70s, giving rise to the PRA, continued to dissipate throughout the 1980s. With organized labor on the defensive, management saw an opportunity.

In 1982, the high-volume plant in Los Angeles was selected to pilot the Postal Service's first multiline OCRs, which could automatically read entire addresses, not just ZIP codes like their single-line predecessors. To celebrate, the Postal Service also coined a new employment category. "Mail Processors," who would monitor these new OCRs, were added to the employment hierarchy two rungs on the payscale below that of MPLSM operators. The APWU filed a grievance, alleging that this constituted unfair labor practices and went against the PRA provision about bargaining over new technology.

Reagan's notoriously management-friendly National Labor Relations Board sided with the Postal Service, allowing the new employment category to stand. It turned out that the consultants had been wrong: it wasn't the technical feat of bringing OCR up to 85 percent

accuracy that made automation economically viable, but rather the Postal Service's technocratic insight about how to redefine work and reduce wages.

This was a clever bit of politicking, but it improved neither the literacy rate of machines nor the USPS's overall service outcomes. Over time, the optimistic narratives of innovation that dominated the 1980s began to wane. Progress was made in OCR, but not quickly enough to keep pace with the steady growth in mail volume. Two congressional reports published in the early nineties captured this sentiment: 1992's "Automation Is Restraining But Not Reducing Costs" and 1995's "Automation Is Taking Longer and Producing Less Than Expected."

Hamstrung by successive waves of neoliberal policymaking, the twin values of service and innovation upon which the Post Office was founded had been rendered incompatible with one another. A regime of unrelenting austerity had motivated and justified a blind faith in the promise of automation—all the while undermining this promise. It was out of this failure that the Remote Encoding Center was born.

Long-Term Temporary

Nairn Higginson was a day-one hire at the Salt Lake City REC when it opened for business in 1994. He had responded to a job ad seeking keyers, which described the position as "long-term temporary" and "strictly transitional." The gig paid more than double the

federal minimum wage, and previous experience as a typist or computer technician was required.

Higginson recounted to me how frequently he and his colleagues had been advised by management over the years that their days were soon to be numbered. More than a quarter-century into his tenure at the USPS, he is now the REC's Manager of Operations. Much remains the same since his days as a keyer, save for one notable difference: the hards have gotten considerably harder. "Those last few percentage points," another veteran keyer hired in the mid-nineties noted, alluding to lingering OCR error rates, "take years and years and years."

In the early 1990s, as a last-ditch effort aimed at bolstering the still imperfect machine-reader systems it had so heavily invested in over the decade prior, the USPS began to experiment with new "remote video encoding" technology. Rather than require on-site clerks to deal with each piece of mail rejected by the machine-readers, "remote encoding" provided an off-site, human backstop to augment the OCR systems. If successful, this fix would serve to prop up automated sorting operations until OCR technology had improved enough to finally make the remaining human cogs in the mailstream redundant, once and for all.

The USPS initially opted for part-time subcontractors located outside of expensive metropolitan hubs, rather than full-time career postal workers. Twenty-five of these remote encoding facilities were launched by private firms in 1992, with another two hundred slated to open in the following years. But the APWU intervened, claiming that this public-private subcontracting arrangement was a breach of the union's collective bargaining agreement. This time, the Clinton-appointed National Labor Relations Board ruled in the union's favor, bringing all remote encoding operations back in-house.

Salt Lake City was the inaugural outpost, and additional RECs in upstate New York, suburban Arkansas, and formerly industrialized East Pittsburgh soon joined the ranks. Many of the subcontracted workers hired and trained by private firms were recruited to join the unionized workforce of USPS keyers. On the eve of the twenty-first century, there were 30,000 keyers at 55 RECs, keying 25 billion images per year. RECs had been conceived under the pretense of imminent redundancy, but had proven surprisingly resilient.

When Higginson started as a keyer, each REC handled the nonmachinable parcels for only its regional processing plant. (Higginson once keyed a letter addressed to himself, sent by a friend with particularly inscrutable handwriting.) Today, however, the first of its kind is the last still standing: all mail that stumps the OCRs in every state and every territory flows through Utah. The barrios of Puerto Rico send a disproportionate share,

several keyers told me, speculating that Postal Service OCRs had not been designed to account for the territory's unique street address schema.

The long-presaged closures and consolidations had finally begun in the first decade of the 2000s. But the new century mirrored the old: tech had improved since the Postal Reorganization Act era, to be sure, but changes in the political climate were more decisive. The Postal Accountability and Enhancement Act, signed by President Bush in 2006, renewed and deepened the reign of postal austerity instituted by Nixon and Reagan.

Only ever a “temporary fix,” RECs were first on the chopping block. By 2007, all but eight were shuttered, and 80 percent of keyers had been laid off. Closures continued, and some seasoned keyers relocated more than 1,000 miles to continue keying at the final two remaining facilities—now deemed MegaRECs—in Salt Lake City and Wichita. The two, however, became one in 2014, when Wichita was decommissioned, and all national remote encoding operations were consolidated in Utah.

Today, only about one-third of the REC's 1,500 beige cubicles are occupied at any given time, as budgets have continued to tighten and postal OCR has improved to read more than 99 percent of letters and about 85 percent of packages. As the overall stock of illegible mail continues to shrink, it follows that each remaining item sent to the REC is progressively more degraded and harder to read. “The quality of the images we get sent,”

Higginson reported, “is getting worse and worse.” Incremental steps forward in the tech make for incremental steps backward at the REC.

Very Hard Tasks

The story of the REC—and its uncertain future—is a parable for what happens when a robust public service is systematically hollowed out by the dagger of neoliberalism. For the Postal Service, this dagger has often been hidden within the cloak of technological solutionism. But the fully automated mailstream, ever over-promised and under-delivered, may finally be materializing. In late 2019, the USPS announced a partnership with NVIDIA, the leading producer of the powerful computer chips that have catalyzed recent breakthroughs in artificial intelligence.²¹⁰

According to the press release, sorting machinery in more than two hundred USPS facilities will soon be enhanced with AI models designed by NVIDIA. The company’s image recognition systems have advanced so significantly in the past decade that they can now reliably distinguish a white wolf from a species of large white dog. Reading the mail, one might assume, should be a no-brainer. Indeed, OCR error rates have plummeted

²¹⁰ Brown, Ken. “NVIDIA Provides U.S. Postal Service AI Technology to Improve Delivery Service.” NVIDIA Newsroom, November 5, 2019. <http://nvidianews.nvidia.com/news/nvidia-provides-u-s-postal-service-ai-technology-to-improve-delivery-service>.

thanks to classification models like the ones developed by NVIDIA, and character recognition is now sometimes referred to by researchers as a “solved” problem. Perhaps the notoriously stubborn “last few percentage points”—the ones that have kept the Utah keyers busy for far longer than anyone expected—have finally met their match.

Curiously, however, the origins of these very same AI systems now installed in mail processing plants—known as “deep” neural networks—can be traced back to the USPS. In the late 1980s, a young researcher at AT&T named Yann LeCun began to experiment with neural networks using a dataset provided by the Postal Service.²¹¹ It contained about 9,000 images of individual digits, culled from handwritten ZIP codes. To this day, a modified version of the dataset is ubiquitous in computer science curriculums, serving as a benchmark for handwriting recognition systems. LeCun, now the head of AI at Facebook, expressed his gratitude to the Postal Service in an early published technical paper, citing the “very hard tasks” performed by USPS’s engineering department in preparing this dataset for use by AT&T researchers.

²¹¹ See LeCun et al., “Backpropagation Applied to Handwritten Zip Code Recognition,” *Neural Computation* 1, no. 4 (1989): 541–51 and “Handwritten Digit Recognition with a Back-Propagation Network,” *Advances in Neural Information Processing Systems (NIPS 1989)* (Denver, Colo., 1990), <http://nyuscholars.nyu.edu/en/publications/handwritten-digit-recognition-with-a-back-propagation-network-2>.

Just a few years before LeCun was given access to the formative ZIP code dataset, the USPS had been forced to acquiesce to AT&T on another front. Despite message volume increasing by a factor of ten in its first two years, the USPS's ambitious E-COM program—a proto-email system—was discontinued after pushback from the private telecommunications industry.²¹² AT&T led the charge, claiming that it was unfair that the “Post Office is being encouraged to provide a kind of service” that “private industry is able to do.” That the Post Office could afford to invest in innovative and promising new technology, even when it was unprofitable, was an outrageous notion. Reagan's Federal Communications Commission agreed, bringing to an untimely demise what may well have amounted to a publicly owned, state-of-the-art digital communications infrastructure.

The handwriting dataset and the “very hard tasks” required to produce it, as well as the preemptive gutting of E-COM, are just two particularly salient examples of an ongoing transfer of postal resources—both intellectual and infrastructural—from the public to the private sector. But providing a public subsidy for private enterprise was nothing new for the Postal Service. Long before Nixon and Reagan, the Post Office was forced to kowtow to the whims of commercial mailers, offering cut-rate postage and special delivery options. And more recently, the USPS expanded its delivery window to include Sundays for the

²¹² Ellis, Ryan. “The Premature Death of Electronic Mail: The United States Postal Service's E-COM Program, 1978-1985.” *International Journal of Communication* 7, no. 0 (August 30, 2013): 19.

first time in its history—except Sundays were reserved exclusively for Amazon packages, which provided a much-needed revenue stream for the fiscally precarious Postal Service. “We deliver Amazon packages until we drop dead,” read the headline of a 2018 USPS letter-carrier tell-all.²¹³ Amazon, in effect, built its empire off the back of public postal infrastructure, all the while scaling up its own massive logistics operation, staffed not by public sector employees, but contractors and subcontractors.

Decades of austerity have driven the Postal Service into a state of submission, depending for its continued sustenance on the goodwill of the same private companies it helped get off the ground. Despite this, the USPS continues to serve as a keystone of democracy. Last fall, it executed an unprecedented vote-by-mail operation under the tutelage of a hostile Postmaster General. It continues to serve as a vital conduit of information for the more than 20 million Americans without internet access.

As postal history demonstrates, time and again, the work of overcoming hard problems like these is rarely just a technical achievement. After all, universal mail service means that the mail must be delivered, hardness notwithstanding, whether by humans, machines, or some combination thereof. And until the fully automated mailstream becomes a reality,

²¹³ O’Connor, Brendan. “Confessions of a U.S. Postal Worker: ‘We Deliver Amazon Packages until We Drop Dead.’” *GEN* (blog), August 30, 2019. <https://gen.medium.com/confessions-of-a-u-s-postal-worker-we-deliver-amazon-packages-until-we-drop-dead-a6e96f125126>.

some of the hardest problems facing the Postal Service will be solved by keyers at the Salt Lake City REC, five million times per day, one billion times per year.

Conclusion

As much ballyhoo as there has been about the recent breakthroughs in AI development, many of the practical implementations of machine learning have been cause for concern for a growing set of critics, pundits, and policymakers.²¹⁴ The digital media theorist Matthew Kirschenbaum has written at length about the apparent threat to human creativity posed by machine learning and artificial intelligence. Recently, he published a portentous account of the impending “textocalypse,” arguing that machine-generated text would soon usurp human-produced texts, following the popularization of generative AI platforms like ChatGPT.²¹⁵ Ultimately, Kirschenbaum raises the alarm about “a crisis of never-ending spam,” which he argues will catalyze an epistemological breakdown where misinformation runs amok, and the demand for textual works – creative, journalistic, and academic alike – created by humans utterly disintegrates. Generative text programs like

²¹⁴ Lieu, Ted. “I’m a Congressman Who Codes. A.I. Freaks Me Out.” *The New York Times*, January 23, 2023, sec. Opinion. <https://www.nytimes.com/2023/01/23/opinion/ted-lieu-ai-chatgpt-congress.html>.
McQuillan, Dan. *Resisting AI: An Anti-Fascist Approach to Artificial Intelligence*. First edition. Bristol: Bristol University Press, 2022.

²¹⁵ Kirschenbaum, Matthew. “Prepare for the Textpocalypse.” *The Atlantic*, March 8, 2023. <https://www.theatlantic.com/technology/archive/2023/03/ai-chatgpt-writing-language-models/673318/>.

ChatGPT, along with image-based analogues like Midjourney, DALL-E, Stable Diffusion and others, indeed seem to be annexing ever greater swathes of the digital landscape.²¹⁶

However, perhaps Kirschenbaum's apprehension is misplaced, and even bordering on ahistorical.²¹⁷ The looming glut of generative media may well have the inverse effect – entrenching the value of human-produced work in spite of the seemingly shrinking gap between human and machine faculties of judgment, interpretation, and creativity. The history of CAPTCHA, to wit, provides a pertinent counterexample, whereby the ostensibly clearcut ontological distinction between humans and machines proves to be a deeply fraught entanglement.²¹⁸ As I have written elsewhere, the genealogy of CAPTCHA

²¹⁶ For a range of recent coverage, including ChatGPT being experimented with in medical, legal, educational, and entertainment sectors, as well as further integration into existing Big Tech platforms, see: Brodwin, Erin. “Doximity’s Beta Tool, DocsGPT, Uses Health Care-Specific Language to Speed Workflow.” *Axios* (blog), February 13, 2023. Journal, A. B. A. “Meet Harvey, BigLaw Firm’s Artificial Intelligence Platform Based on ChatGPT.” *ABA Journal*. Accessed March 9, 2023. Leswing, Kif. “ChatGPT Is Being Used to Automatically Write Emails: Microsoft, Salesforce and TikTok Creators Are Hopping on the Trend.” *CNBC*, March 8, 2023. Wilkes, Emma. “David Guetta Says ‘the Future of Music Is in AI’ after Eminem Deepfake Vocal Stunt.” *NME* (blog), February 14, 2023. Quizlet. “Introducing Q-Chat, the World’s First AI Tutor Built with OpenAI’s ChatGPT.” Accessed March 9, 2023.

²¹⁷ Joseph Bernstein has produced arguably one of the sharpest historicizations of the ongoing panic around misinformation. Bernstein, Joseph. “Bad News: Selling the Story of Disinformation.” *Harper’s Magazine*, August 9, 2021. <https://harpers.org/archive/2021/09/bad-news-selling-the-story-of-disinformation/>.

²¹⁸ For more on the relationship between CAPTCHA and misinformation, and the processes of reification associated with “bots,” see: Justie, Brian. “Bot or Not.” *Real Life*, September 21, 2020. <https://reallifemag.com/bot-or-not/>.

and reCAPTCHA powerfully demonstrates that the specter of “bots,” and misinformation more generally, can be easily construed as just another vector of social reification that fuels further overreach of data capture and valorization across the internet. As such, emergent forms of AI, like those analyzed in Chapter 2, exemplify this tendency and merit scrutiny not just as technical curios, but as social agents embedded within contentious political-economic systems.

Further, Chapters 1 and 3 provide direct examples of the forms of labor that are repeatedly revealed to undergird the operations of these systems – the former in a distributed and casualized context steadily expanding its reach to encompass nearly all users of the increasingly privatized contemporary internet, and the latter in a high-stakes professional setting subjected to the downward pressures of public sector austerity. In both cases, vast systems of digital infrastructure depend upon the continued access to a pool of humans performing micro-tasks. These tasks – regardless of whether or not they are understood to be *work* per se – will arguably prove increasingly valuable in the near-term future, even if they tend toward banality, gradually approximating something that feels more and more robotic and mindless.

Asymptotic labor, therefore, is a concept with both diagnostic and speculative appeal. It can be used to diagnose the ways that Big Tech has upended both the formal and informal political economy, as well as provide a template for speculating about the future of data-

intensive – which is to say, *labor-intensive* – information-processing systems and the implications of their increasing pervasiveness. Ultimately, it serves to reveal a perverse irony attendant to the contemporary technological moment, wherein the tasks that ultimately prove one’s humanity in the face of ever-encroaching digital automata demand that we debase ourselves, bit by bit, as our actions become more machine-like and inhuman.

Bibliography

@dril_gpt2. Twitter Post. January 27, 2020, 3:03 a.m.

https://twitter.com/dril_gpt2/status/1221750499688550406.

“Reevaluating the Universal Service Obligation.” Office of Inspector General, United States Postal Service, May 6, 2020.

<https://www.uspsoig.gov/sites/default/files/reports/2023-01/RISC-WP-20-004.pdf>.

Abraham, Tara. *Rebel Genius: Warren S. McCulloch’s Transdisciplinary Life in Science*. 1st edition. Cambridge, Massachusetts: The MIT Press, 2016.

Agre, Philip E. “Conceptions of the User in Computer Systems Design.” In *The Social and Interactional Dimensions of Human-Computer Interfaces*, 67–106. USA: Cambridge University Press, 1995.

Agre, Philip E. “Surveillance and Capture: Two Models of Privacy.” *The Information Society* 10, no. 2 (April 1994): 101–27.

<https://doi.org/10.1080/01972243.1994.9960162>.

Agre, Philip E. *Computation and Human Experience*. Illustrated edition. Cambridge ; New York: Cambridge University Press, 1997.

Ahn, Luis von, and Laura Dabbish. “Designing Games with a Purpose.”

Communications of the ACM 51, no. 8 (August 1, 2008): 57.

<https://doi.org/10.1145/1378704.1378719>.

Ahn, Luis von, and Laura Dabbish. “Labeling Images with a Computer Game.” In *Proceedings of the 2004 Conference on Human Factors in Computing Systems - CHI ’04*, 319–26. Vienna, Austria: ACM Press, 2004.

<https://doi.org/10.1145/985692.985733>.

Ahn, Luis von, Benjamin Maurer, Colin McMillen, David Abraham, and Manuel Blum. “ReCAPTCHA: Human-Based Character Recognition via Web Security

- Measures.” *Science* 321, no. 5895 (September 12, 2008): 1465–68.
<https://doi.org/10.1126/science.1160379>.
- Ahn, Luis von, Manuel Blum, and John Langford. “CMU-CS-02-117 (‘Telling Humans and Computers Apart (Automatically), or How Lazy Cryptographers Do AI’).” Technical Report. School of Computer Science, Carnegie Mellon University, February 2002. <http://reports-archive.adm.cs.cmu.edu/anon/2002/abstracts/02-117.html>.
- Ahn, Luis von, Manuel Blum, and John Langford. “Telling Humans and Computers Apart Automatically.” *Communications of the ACM* 47, no. 2 (February 1, 2004): 56–60. <https://doi.org/10.1145/966389.966390>.
- Ahn, Luis von, Manuel Blum, Nicholas Hopper, and John Langford. “Captcha: Telling Humans and Computers Apart Automatically.” In *Proceedings of Eurocrypt*, 294–311. Warsaw, Poland: Springer, 2003. https://doi.org/10.1007/3-540-39200-9_18.
- Ahn, Luis von, Manuel Blum, Nicholas J. Hopper, and John Langford. “CAPTCHA: Using Hard AI Problems for Security.” In *Advances in Cryptology — EUROCRYPT 2003*, edited by Eli Biham, 2656:294–311. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003. https://doi.org/10.1007/3-540-39200-9_18.
- Ahn, Luis von, Mihir Kedia, and Manuel Blum. “Verbosity: A Game for Collecting Common-Sense Facts.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI ’06*, 75. Montreal, Canada: ACM Press, 2006. <https://doi.org/10.1145/1124772.1124784>.
- Ahn, Luis von, Ruoran Liu, and Manuel Blum. “Peekaboom: A Game for Locating Objects in Images.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI ’06*, 55. Montreal, Canada: ACM Press, 2006.
<https://doi.org/10.1145/1124772.1124782>.
- Ahn, Luis von. “Human Computation.” Dissertation, Carnegie Mellon University, 2005.

- Ahn, Luis von. “Teaching Computers to Read: Google Acquires ReCAPTCHA.” *Official Google Blog* (blog), September 16, 2009.
<https://googleblog.blogspot.com/2009/09/teaching-computers-to-read-google.html>.
- Alammar, Jay. “The Illustrated Transformer,” June 27, 2018.
<https://jalammar.github.io/illustrated-transformer/>.
- Alimahomed-Wilson, Jake. “The Amazonification of Logistics: E-Commerce, Labor, and Exploitation in the Last Mile.” In *The Cost of Free Shipping*, edited by Jake Alimahomed-Wilson and Ellen Reese, 69–84. Amazon in the Global Economy. Pluto Press, 2020. <https://doi.org/10.2307/j.ctv16zjhj.11>.
- Altenried, Moritz. “On the Last Mile: Logistical Urbanism and the Transformation of Labour.” *Work Organisation, Labour & Globalisation* 13 (April 1, 2019): 114–29.
<https://doi.org/10.13169/workorgalaboglob.13.1.0114>.
- Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner. “Machine Bias.” *ProPublica*, May 23, 2016. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
- Baird, Henry S., Allison L. Coates, and Richard J. Fateman. “Pessimist Print: A Reverse Turing Test.” In *Proceedings of the 6th International Conference on Document Analysis and Recognition*, 1154–58. Seattle, WA: IEEE Computer Society, 2001.
- Baird, Henry S., and Kris Popat. “Human Interactive Proofs and Document Image Analysis.” In *Document Analysis Systems V*, edited by Daniel Lopresti, Jianying Hu, and Ramanujan Kashi, 2423:507–18. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002. https://doi.org/10.1007/3-540-45869-7_54.
- Baker, Paul, and Amanda Potts. “‘Why Do White People Have Thin Lips?’ Google and the Perpetuation of Stereotypes via Auto-Complete Search Forms.” *Critical Discourse Studies* 10, no. 2 (May 1, 2013): 187–204.
<https://doi.org/10.1080/17405904.2012.744320>.

- Barthes, Roland. *Camera Lucida: Reflections on Photography*. Translated by Richard Howard. Reprint edition. New York: Hill and Wang, 2010.
- Baxter, Vern K. *Labor and Politics in the U.S. Postal Service*. Plenum Studies in Work and Industry. Springer, 2013.
- BBC News. “PC Stripper Helps Spam to Spread,” October 30, 2007.
<http://news.bbc.co.uk/2/hi/technology/7067962.stm>.
- Bender, Emily, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” *FAccT ‘21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, March 2021, 610–23.
- Benjamin, Walter. *The Work of Art in the Age of Its Technological Reproducibility, and Other Writings on Media*. Harvard University Press, 2008.
- Bernstein, Joseph. “Bad News: Selling the Story of Disinformation.” *Harper’s Magazine*, August 9, 2021. <https://harpers.org/archive/2021/09/bad-news-selling-the-story-of-disinformation/>.
- Blok, Aad, Greg Downey, and Greg Downey. *Uncovering Labour in Information Revolutions, 1750-2000: Volume 11*. Cambridge University Press, 2003.
- Bourdieu, Pierre. *Distinction: A Social Critique of the Judgement of Taste*. Harvard University Press, 1984.
- boyd, danah. “Did Media Literacy Backfire?” *Points (Data & Society)* (blog), March 16, 2018. <https://points.datasociety.net/did-media-literacy-backfire-7418c084d88d>.
- Braverman, Harry. *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. 25th Anniversary ed. edition. New York: Monthly Review Press, 1998.
- Breiman, Leo. “Statistical Modeling: The Two Cultures (with Comments and a Rejoinder by the Author).” *Statistical Science* 16, no. 3 (August 2001): 199–231. <https://doi.org/10.1214/ss/1009213726>.

- Brodwin, Erin. “Doximity’s Beta Tool, DocsGPT, Uses Health Care-Specific Language to Speed Workflow.” Axios (blog), February 13, 2023.
<https://www.axios.com/pro/health-tech-deals/2023/02/13/chatgpt-healthcare-uses-examples-doximity>.
- Brown, Ken. “NVIDIA Provides U.S. Postal Service AI Technology to Improve Delivery Service.” NVIDIA Newsroom, November 5, 2019.
<http://nvidianews.nvidia.com/news/nvidia-provides-u-s-postal-service-ai-technology-to-improve-delivery-service>.
- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. “Language Models Are Few-Shot Learners.” arXiv, July 22, 2020. <https://doi.org/10.48550/arXiv.2005.14165>.
- Brunton, Finn. *SPAM: A Shadow History of the Internet*. Infrastructures. MIT Press, 2013.
- Buckner, Cameron, and James Garson. “Connectionism and Post-Connectionist Models.” In *The Routledge Handbook of the Computational Mind*. Routledge, 2018.
- Buckner, Cameron. “Empiricism without Magic: Transformational Abstraction in Deep Convolutional Neural Networks.” *Synthese* 195, no. 12 (December 1, 2018): 5339–72. <https://doi.org/10.1007/s11229-018-01949-1>.
- Bunz, Mercedes. “The Calculation of Meaning: On the Misunderstanding of New Artificial Intelligence as Culture.” *Culture, Theory and Critique* 60, no. 3–4 (October 2, 2019): 264–78. <https://doi.org/10.1080/14735784.2019.1667255>.
- Burrell, Jenna. “Thinking Relationally about Digital Inequality in Rural Regions of the U.S.” *First Monday* 23, no. 6 (June 1, 2018).
<https://doi.org/10.5210/fm.v23i6.8376>.
- Cameron Buckner and James Garson, “Connectionism,” Stanford Encyclopedia of Philosophy Archive, 16 Aug. 2019,
<http://plato.stanford.edu/archives/fall2019/entries/connectionism>.

- Cardon, Dominique, Jean-Philippe Cointet, and Antoine Mazières. “Neurons Spike Back (La Revanche Des Neurones: L’invention Des Machines Inductives et La Controverse de l’intelligence Artificielle).” *Réseaux* n° 211, no. 5 (2018): 173. <https://doi.org/10.3917/res.211.0173>.
- Chen, Angela. “Inmates in Finland Are Training AI as Part of Prison Labor.” The Verge (blog), March 28, 2019. <https://www.theverge.com/2019/3/28/18285572/prison-labor-finland-artificial-intelligence-data-tagging-vainu>.
- Chew, Monica, and Henry S. Baird. “BaffleText: A Human Interactive Proof.” In *Proceedings of Document Recognition and Retrieval X*, edited by Tapas Kanungo, Elisa H. Barney Smith, Jianying Hu, and Paul B. Kantor, 5010:305–16. Santa Clara, CA, 2003. <https://doi.org/10.1117/12.479682>.
- Chew, Monica, and J. D. Tygar. “Image Recognition CAPTCHAs.” In *Information Security*, edited by Kan Zhang and Yuliang Zheng, 3225:268–79. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004. https://doi.org/10.1007/978-3-540-30144-8_23.
- Churchland, Patricia S., and Terrence J. Sejnowski. *The Computational Brain*. Anniversary ed. edition. Cambridge, MA: Bradford Books, 2016.
- Coates, A.L., H.S. Baird, and R.J. Faternan. “Pessimist Print: A Reverse Turing Test.” In *Proceedings of Sixth International Conference on Document Analysis and Recognition*, 1154–58, 2001. <https://doi.org/10.1109/ICDAR.2001.953966>.
- Cockburn, Cynthia. *Machinery of Dominance: Women, Men and Technical Know-How*. London ; Dover, N.H: Pluto Press, 1987.
- Cope, Jonathan. “Neoliberalism and Library & Information Science: Using Karl Polanyi’s Fictitious Commodity as an Alternative to Neoliberal Conceptions of Information.” *Publications and Research*, January 1, 2014. https://academicworks.cuny.edu/si_pubs/4.

- Cordell, Ryan. “‘Q i-Jtb the Raven’: Taking Dirty OCR Seriously.” *Book History* 20, no. 1 (2017): 188–225. <https://doi.org/10.1353/bh.2017.0006>.
- Cowan, Ruth Schwartz. *More Work For Mother*. Basic Books, 1983.
- Crawford, Kate, and Trevor Paglen. “Excavating AI: The Politics of Training Sets for Machine Learning.” *Excavating AI*, September 19, 2019. www.excavating.ai.
- Danchev, Dancho. “Inside India’s CAPTCHA Solving Economy.” *ZDNet* (blog), August 29, 2008. <https://www.zdnet.com/article/inside-indias-captcha-solving-economy/>.
- Daston, Lorraine. “Calculation and the Division of Labor, 1750-1950.” *Bulletin of the German Historical Institute* 62, no. Spring (2018): 9–30.
- Deacon, Terrence W. “Shannon - Boltzmann — Darwin: Redefining Information (Part II).” *Cognitive Semiotics* 2, no. s1 (March 1, 2008): 169–96. <https://doi.org/10.1515/cogsem.2008.2.spring2008.169>.
- Deacon, Terrence W. *The Symbolic Species: The Co-Evolution of Language and the Brain*. New York, NY: W. W. Norton & Company, 1998.
- Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. “ImageNet: A Large-Scale Hierarchical Image Database.” In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–55, 2009. <https://doi.org/10.1109/CVPR.2009.5206848>.
- Derrida, Jacques. “Structure, Sign, and Play in the Discourse of the Human Sciences.” In *The Structuralist Controversy: The Languages of Criticism and the Sciences of Man*, edited by Richard A. Macksey, 40th anniversary edition., 186–94. Baltimore, Md: Johns Hopkins University Press, 2007.
- Doane, Mary Ann. “Indexicality: Trace and Sign: Introduction.” *Differences* 18, no. 1 (May 1, 2007): 1–6. <https://doi.org/10.1215/10407391-2006-020>.
- Doctorow, Cory. “Solving and Creating Captchas with Free Porn.” *Boing Boing* (blog), May 8, 2004. <https://boingboing.net/2004/01/27/solving-and-creating.html>.

- Dreyfus, Hubert L. *What Computers Can't Do: The Limits of Artificial Intelligence*. Revised, Subsequent edition. New York: HarperCollins, 1978.
- Drucker, Johanna. "Humanities Approaches to Graphical Display." *Digital Humanities Quarterly* 005, no. 1 (March 10, 2011).
- Drucker, Johanna. *Graphesis: Visual Forms of Knowledge Production*. Harvard University Press, 2014.
- Dutton, William H., and Kenneth L. Kraemer. "Automating Bias." *Society* 17, no. 2 (January 1, 1980): 36–41. <https://doi.org/10.1007/BF02700058>.
- Dyer-Witheford, Nick, Atle Mikkola Kjosén, and James Steinhoff. *Inhuman Power: Artificial Intelligence and the Future of Capitalism*. London: Pluto Press, 2019.
- Egele, Manuel, Leyla Bilge, Engin Kirda, and Christopher Kruegel. "CAPTCHA Smuggling: Hijacking Web Browsing Sessions to Create CAPTCHA Farms." *SAC '10: Proceedings of the 2010 ACM Symposium on Applied Computing*, 2010, 1865–70. <http://dx.doi.org/10.1145/1774088.1774483>.
- Ellis, Ryan. "The Premature Death of Electronic Mail: The United States Postal Service's E-COM Program, 1978-1985." *International Journal of Communication* 7, no. 0 (August 30, 2013): 19
- Ellis, Ryan. *Letters, Power Lines, and Other Dangerous Things: The Politics of Infrastructure Security*. MIT Press, 2020
- Engber, Daniel. "Who Made That Captcha?" *The New York Times*, January 17, 2014, sec. Magazine. <https://www.nytimes.com/2014/01/19/magazine/who-made-that-captcha.html>.
- Fairchild, Mary Salome Cutler. "Women in American Libraries." *Library Journal* 29 (1904): 157–62.
- Fazi, M. Beatrice. "Beyond Human: Deep Learning, Explainability and Representation." *Theory, Culture & Society* 38, no. 7–8 (December 1, 2021): 55–77. <https://doi.org/10.1177/0263276420966386>.

- Fisher, Mark. *Capitalist Realism: Is There No Alternative?* Winchester, UK
Washington, USA: Zero Books, 2009.
- Fodor, J., and B. P. McLaughlin. “Connectionism and the Problem of Systematicity:
Why Smolensky’s Solution Doesn’t Work.” *Cognition* 35, no. 2 (May 1990): 183–
204. [https://doi.org/10.1016/0010-0277\(90\)90014-b](https://doi.org/10.1016/0010-0277(90)90014-b).
- Fraser, Nancy. “Behind Marx’s Hidden Abode.” *New Left Review*, no. 86 (April 1, 2014):
55–72.
- Fukushima, Kunihiro. “Neocognitron: A Self-Organizing Neural Network Model for a
Mechanism of Pattern Recognition Unaffected by Shift in Position.” *Biological
Cybernetics* 36, no. 4 (April 1, 1980): 193–202.
<https://doi.org/10.1007/BF00344251>.
- Fuller, Wayne. *The American Mail: Enlarger of the Common Life*. Chicago: University
of Chicago Press, 1972.
- Fussell, Sidney. “Behind Every Robot Is a Human.” *The Atlantic*, April 15, 2019.
[https://www.theatlantic.com/technology/archive/2019/04/amazon-workers-
eavesdrop-amazon-echo-clips/587110/](https://www.theatlantic.com/technology/archive/2019/04/amazon-workers-eavesdrop-amazon-echo-clips/587110/).
- Gal, Susan, and Judith T. Irvine. *Signs of Difference: Language and Ideology in Social
Life*. Cambridge, United Kingdom ; New York, NY: Cambridge University Press,
2019.
- Galloway, Alexander. “The Poverty of Philosophy: Realism and Post-Fordism.” *Critical
Inquiry* 39, no. 2 (2013): 347–66. <https://doi.org/10.1086/668529>.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep Learning*. Illustrated
edition. Cambridge, Massachusetts: The MIT Press, 2016.
- Google Developers. “Choosing the Type of ReCAPTCHA.” Accessed January 30, 2020.
<https://developers.google.com/recaptcha/docs/versions>.
- Google Developers. “Invisible ReCAPTCHA.” Accessed January 30, 2020.
<https://developers.google.com/recaptcha/docs/invisible>.

- GPT-3: Language Models Are Few-Shot Learners (Paper Explained), 2020.
<https://www.youtube.com/watch?v=SY5PvZrJhLE>.
- Gray, Mary L., and Siddharth Suri. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Eamon Dolan/Houghton Mifflin Harcourt, 2019.
- Green, Venus. *Race on the Line: Gender, Labor, and Technology in the Bell System, 1880-1980*. Duke University Press, 2001.
- Gugliotta, Guy. "Captchas Have Us Deciphering Old Text Through Woozy Web Clues." *The New York Times*, March 28, 2011, sec. Science.
<https://www.nytimes.com/2011/03/29/science/29recaptcha.html>.
- Gunning, Tom. "What's the Point of an Index? Or, Faking Photographs." In *Still Moving: Between Cinema and Photography*, edited by Karen Beckman and Jean Ma. Duke University Press, 2008.
- Hacking, Ian. *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge University Press, 1983.
- Halevy, Alon, Peter Norvig, and Fernando Pereira. "The Unreasonable Effectiveness of Data." *IEEE Intelligent Systems* 24, no. 2 (March 2009): 8–12.
<https://doi.org/10.1109/MIS.2009.36>.
- Halpern, Orit. *Beautiful Data: A History of Vision and Reason since 1945*. Illustrated edition. Durham: Duke University Press Books, 2015.
- Harper's New Monthly Magazine*. "New York City Post-Office," November 1871.
- Hawkes, Terence. *Structuralism & Semiotics*. University of California Press, 1977.
- Hayles, N. Katherine. "Can Computers Create Meanings? A Cyber/Bio/Semiotic Perspective." *Critical Inquiry* 46, no. 1 (September 2019): 32–55.
<https://doi.org/10.1086/705303>.
- Hernandez, Danny, and Dario Amodei. "AI and Compute." *OpenAI* (blog), May 16, 2018. <https://openai.com/blog/ai-and-compute/>.

- Hicks, Marie. *Programmed Inequality: How Britain Discarded Women Technologists and Lost Its Edge in Computing*. MIT Press, 2017.
- Hochreiter, Sepp, and Jürgen Schmidhuber. “Long Short-Term Memory.” *Neural Computation* 9, no. 8 (November 15, 1997): 1735–80.
<https://doi.org/10.1162/neco.1997.9.8.1735>.
- Hopper, Nicholas J., and Manuel Blum. “Secure Human Identification Protocols.” In *Advances in Cryptology — ASIACRYPT 2001*, edited by Colin Boyd, 2248:52–66. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001.
https://doi.org/10.1007/3-540-45682-1_4.
- Hopper, Nicholas, John Langford, and Luis von Ahn. “Provably Secure Steganography.” In *CRYPTO 2002: Advances in Cryptology*, 2002. https://doi.org/10.1007/3-540-45708-9_6.
- Hubel, D. H., and T. N. Wiesel. “Receptive Fields, Binocular Interaction and Functional Architecture in the Cat’s Visual Cortex.” *The Journal of Physiology* 160, no. 1 (January 1962): 106-154.2.
- Huet, Ellen. “The Humans Hiding Behind the Chatbots.” *Bloomberg*, April 18, 2016.
<https://www.bloomberg.com/news/articles/2016-04-18/the-humans-hiding-behind-the-chatbots>.
- Human Computation*. Google TechTalk. Mountain View, CA, 2006.
<https://www.youtube.com/watch?v=tx082gDwGcM>.
- Hwang, Tim. “Computational Power and the Social Impact of Artificial Intelligence.” *SSRN Scholarly Paper*. Rochester, NY: Social Science Research Network, March 23, 2018. <https://papers.ssrn.com/abstract=3147971>.
- Irani, Lilly. “Difference and Dependence among Digital Workers: The Case of Amazon Mechanical Turk.” *South Atlantic Quarterly* 114, no. 1 (2015): 225–34.
<https://doi.org/10.1215/00382876-2831665>.

- Irani, Lilly. *Chasing Innovation: Making Entrepreneurial Citizens in Modern India*. Princeton University Press, 2019.
- Jakobson, Roman. “Shifters, Verbal Categories, and the Russian Verb.” In *Russian and Slavic Grammar: Studies 1931-1981*, edited by Linda R. Waugh and Morris Halle. New York: Mouton Publishers, 1984.
- Johnson, George. “To Test a Powerful Computer, Play an Ancient Game.” *The New York Times*, July 29, 1997, sec. Science.
<https://www.nytimes.com/1997/07/29/science/to-test-a-powerful-computer-play-an-ancient-game.html>.
- Journal, A. B. A. “Meet Harvey, BigLaw Firm’s Artificial Intelligence Platform Based on ChatGPT.” *ABA Journal*. Accessed March 9, 2023.
<https://www.abajournal.com/news/article/meet-harvey-biglaw-firms-artificial-intelligence-platform-based-on-chatgpt>.
- Justie, Brian. “Review: Inhuman Power: Artificial Intelligence and the Future of Capitalism.” *Information, Communication & Society* 23, no. 1 (January 2, 2020): 151–54. <https://doi.org/10.1080/1369118X.2019.1651372>.
- Justie, Brian. “Bot or Not.” *Real Life*, September 21, 2020. <https://reallifemag.com/bot-or-not/>.
- Justie, Brian. “Little History of CAPTCHA.” *Internet Histories* 5, no. 1 (January 2, 2021): 30–47. <https://doi.org/10.1080/24701475.2020.1831197>.
- Justie, Brian. “The Nonmachinables.” *Logic Magazine*, 2021.
<https://logicmag.io/distribution/the-nonmachinables/>.
- Karpathy, Andrej. “The Unreasonable Effectiveness of Recurrent Neural Networks.” Andrej Karpathy Blog, May 21, 2015. <http://karpathy.github.io/2015/05/21/rnn-effectiveness/>.

- Keilty, Patrick. "Tedious: Feminized Labor in Machine-Readable Cataloging." *Feminist Media Studies* 18, no. 2 (March 4, 2018): 191–204.
<https://doi.org/10.1080/14680777.2017.1308410>.
- Kelleher, John D. *Deep Learning*. Illustrated edition. Cambridge, Massachusetts: The MIT Press, 2019.
- Kellogg, Sam P. "The Mountain in the Machine: Optimization and the Landscapes of Machine Learning." *Culture Machine* 20 (2021).
- Khan, Lina. "Amazon's Antitrust Paradox." *Yale Law Journal* 126 (January 1, 2017): 710.
- Kirschenbaum, Matthew. "Prepare for the Textpocalypse." *The Atlantic*, March 8, 2023.
<https://www.theatlantic.com/technology/archive/2023/03/ai-chatgpt-writing-language-models/673318/>.
- Kockelman, Paul. "The Epistemic and Performative Dynamics of Machine Learning Praxis." *Signs and Society* 8, no. 2 (March 2020): 319–55.
<https://doi.org/10.1086/708249>.
- Koopman, Colin. *How We Became Our Data: A Genealogy of the Informational Person*. Chicago, IL: University of Chicago Press, 2019.
- Korpela, E., D. Werthimer, D. Anderson, J. Cobb, and M. Leboisky. "SETI@home-Massively Distributed Computing for SETI." *Computing in Science Engineering* 3, no. 1 (January 2001): 78–83. <https://doi.org/10.1109/5992.895191>.
- Krauss, Rosalind. "Notes on the Index: Seventies Art in America." *October* 3 (1977): 68–81. <https://doi.org/10.2307/778437>.
- Kris Paulsen, "The Index and the Interface," *Representations*, no. 122 (Spring 2013): 83–109
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "ImageNet Classification with Deep Convolutional Neural Networks." *Communications of the ACM* 60, no. 6 (May 24, 2017): 84–90. <https://doi.org/10.1145/3065386>.

- Lazer, David, Matthew Baum, and Nicco Mele. “Combating Fake News: An Agenda for Research and Action.” Harvard Kennedy School, May 2017.
<https://shorensteincenter.org/combating-fake-news-agenda-for-research/>.
- Lecun, Yann, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L.D. Jackel. “Backpropagation Applied to Handwritten Zip Code Recognition.” *Neural Computation* 1, no. 4 (1989): 541–51.
- Lecun, Yann, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L.D. Jackel. “Handwritten Digit Recognition with a Back-Propagation Network.” Edited by David Touretzky. *Advances in Neural Information Processing Systems (NIPS 1989)*, Denver, CO 2 (1990).
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. “Deep Learning.” *Nature* 521, no. 7553 (May 2015): 436–44. <https://doi.org/10.1038/nature14539>.
- Lecun, Yann. “A Theoretical Framework for Back-Propagation.” Edited by D. Touretzky, G. Hinton, and T. Sejnowski. *Proceedings of the 1988 Connectionist Models Summer School*, CMU, Pittsburg, PA, 1988, 21–28.
- Leswing, Kif. “ChatGPT Is Being Used to Automatically Write Emails: Microsoft, Salesforce and TikTok Creators Are Hopping on the Trend.” *CNBC*, March 8, 2023. <https://www.cnbc.com/2023/03/08/chatgpt-is-being-used-to-write-emails-big-companies-are-embracing-it.html>.
- Lieu, Ted. “I’m a Congressman Who Codes. A.I. Freaks Me Out.” *The New York Times*, January 23, 2023, sec. Opinion.
<https://www.nytimes.com/2023/01/23/opinion/ted-lieu-ai-chatgpt-congress.html>.
- Light, Jennifer S. “When Computers Were Women.” *Technology and Culture* 40, no. 3 (1999): 455–83.
- Lillibridge, Mark D., Martin Abadi, Krishna Bharat, and Andrei Z. Broder. *Method for selectively restricting access to computer systems*. United States US6195698B1,

- filed April 13, 1998, and issued February 27, 2001.
<https://patents.google.com/patent/US6195698B1/en>.
- Mai, Jens-Erik. “Semiotics and Indexing: An Analysis of the Subject Indexing Process.” *Journal of Documentation* 57, no. 5 (January 1, 2001): 591–622.
<https://doi.org/10.1108/EUM0000000007095>.
- Massumi, Brian. “On the Superiority of the Analog.” In *Parables for the Virtual: Movement, Affect, Sensation*. Duke University Press, 2002.
- May, Matt. “Inaccessibility of CAPTCHA: Alternatives to Visual Turing Tests on the Web.” *Accessible Platform Architectures*. WAI Protocols and Formats Working Group, 2019 2003. <https://www.w3.org/TR/turingtest/#the-accessibility-challenge>.
- McCulloch, Warren S., and Walter Pitts. “A Logical Calculus of the Ideas Immanent in Nervous Activity.” *The Bulletin of Mathematical Biophysics* 5, no. 4 (December 1, 1943): 115–33. <https://doi.org/10.1007/BF02478259>.
- McKenzie Wark, *Capital Is Dead: Is This Something Worse?* (New York, 2019), pp. 45, 46.
- McQuillan, Dan. *Resisting AI: An Anti-Fascist Approach to Artificial Intelligence*. First edition. Bristol: Bristol University Press, 2022.
- Merritt, Rick. “Sharpening Its Edge: U.S. Postal Service Opens AI Apps on Edge Network.” NVIDIA Blog, May 6, 2021.
<https://blogs.nvidia.com/blog/2021/05/06/edge-ai-usps/>.
- Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. “Distributed Representations of Words and Phrases and Their Compositionality.” In *Advances in Neural Information Processing Systems*, Vol. 26. Curran Associates, Inc., 2013.
<https://proceedings.neurips.cc/paper/2013/hash/9aa42b31882ec039965f3c4923ce901b-Abstract.html>.

- Minsky, Marvin, and Seymour Papert. *Perceptrons; an Introduction to Computational Geometry*. MIT Press, 1969.
- Mitchell, W. J. T. *Iconology: Image, Text, Ideology*. Chicago, IL: University of Chicago Press, 1987.
<https://press.uchicago.edu/ucp/books/book/chicago/I/bo5956563.html>.
- Mitchell, William J. *The Reconfigured Eye: Visual Truth in the Post-Photographic Era*. 1st edition. Cambridge, Mass. London: The MIT Press, 1992.
- Mori, G., and J. Malik. “Recognizing Objects in Adversarial Clutter: Breaking a Visual CAPTCHA.” In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2003. Proceedings., I-134-I-141. Madison, WI, USA: IEEE Comput. Soc, 2003. <https://doi.org/10.1109/CVPR.2003.1211347>.
- Motoyama, Marti, Kirill Levchenko, Chris Kanich, Damon McCoy, Geoffrey M Voelker, and Stefan Savage. “Re: CAPTCHAs – Understanding CAPTCHA-Solving Services in an Economic Context.” *Proceedings of the USENIX Security*, 2010.
- Motoyama, Marti. “Understanding the Role of Outsourced Labor in Web Service Abuse.” Dissertation, University of California, San Diego, 2011.
<https://escholarship.org/content/qt87s441sw/qt87s441sw>.
- Nakassis, Constantine V. “Indexicality’s Ambivalent Ground.” *Signs and Society* 6, no. 1 (January 2018): 281–304. <https://doi.org/10.1086/694753>.
- Naor, Moni. “Verification of a Human in the Loop or Identifications via the Turing Test.” Preliminary Draft. Weizmann Institute of Science, September 13, 1996.
- Naor, Moni, and Adi Shamir. “Visual Cryptography.” In *Advances in Cryptology — EUROCRYPT’94*, edited by Alfredo De Santis, 1–12. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 1995. <https://doi.org/10.1007/BFb0053419>.
- Newell, Allen, and Herbert A. Simon. “Computer Science as Empirical Inquiry: Symbols and Search.” *Communications of the ACM* 19 (March 1976).

- Nicholson, Karen P. "The McDonaldization of Academic Libraries and the Values of Transformational Change." *College & Research Libraries* 76, no. 3 (March 1, 2015): 328–38. <https://doi.org/10.5860/crl.76.3.328>.
- Nilsson, Nils J. *Artificial Intelligence: A New Synthesis*. 1st edition. San Francisco, Calif: Morgan Kaufmann Publishers, Inc., 1998.
- Noble, Safiya Umoja. *Algorithms of Oppression: How Search Engines Reinforce Racism*. Illustrated edition. New York: NYU Press, 2018.
- Nöth, Winfried. *Semiotic Foundations of Iconicity in Language and Literature*. Edited by Olga Fischer and Max Nänny. The Motivated Sign. Iconicity in Language and Literature. Philadelphia: John Benjamins Publishing Company, 2001. <https://doi.org/10.1075/ill.2.05not>.
- O'Connor, Brendan. "Confessions of a U.S. Postal Worker: 'We Deliver Amazon Packages until We Drop Dead.'" *GEN*, August 30, 2019. <https://gen.medium.com/confessions-of-a-u-s-postal-worker-we-deliver-amazon-packages-until-we-drop-dead-a6e96f125126>.
- Palo Alto Research Center. "Welcome to PARC's CAPTCHAs." CAPTCHA - History, February 28, 2003. <https://web.archive.org/web/20030407203015/http://www2.parc.com/istl/projects/captcha/history.htm>.
- Parikka, Jussi, and Tony D. Sampson. *The Spam Book: On Viruses, Porn, and Other Anomalies from the Dark Side of Digital Culture*. Hampton Press, 2009.
- Parisi, Luciana. "Reprogramming Decisionism." *E-Flux* 85 (October 2017). www.e-flux.com/journal/85/155472/reprogramming-decisionism/.
- Parmentier, Richard J. *Signs in Society: Studies in Semiotic Anthropology*. Bloomington: Indiana University Press, 1994.
- Pasquale, Frank. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press, 2016.

- Peirce, Charles S. *Charles S. Peirce: Selected Writings*. Edited by Philip P. Wiener. Revised ed. edition. New York: Dover, 1966.
- Peirce, Charles S. *Collected Papers of Charles Sanders Peirce*. Edited by Charles Hartshorne, Paul Weiss, and Arthur W. Burks. Cambridge, Massachusetts: The Belknap Press of Harvard University Press, 1931.
- Peirce, Charles Sanders. *Elements of Logic [Collected Papers of Charles Sanders Peirce, Volume III]*. Edited by Charles Hartshorne and Paul Weiss. Cambridge, MA: Harvard University Press, 1932.
- Peterson, G.D., and R.D. Chamberlain. “Stealing Cycles: Can We Get Along?” In *Proceedings of the Twenty-Eighth Annual Hawaii International Conference on System Sciences*, 2:422–31 vol.2, 1995.
<https://doi.org/10.1109/HICSS.1995.375516>.
- Pinch, Trevor J., and Wiebe E. Bijker. “The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other.” *Social Studies of Science* 14, no. 3 (1984): 399–441.
- Post Office Act of 1792,
<https://collections.si.edu/search/record/ark:/65665/hm8c0ae955a6e4f445e93c830f692e2c1b4>.
- Quizlet. “Introducing Q-Chat, the World’s First AI Tutor Built with OpenAI’s ChatGPT.” Accessed March 9, 2023. https://quizlet.com/blog/meet-q-chat?utm_source=substack&utm_medium=email.
- Radford, Alec, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. “Improving Language Understanding by Generative Pre-Training,” 2018.
- Ray Kurzweil on “I’ve Got a Secret,” 2007.
<https://www.youtube.com/watch?v=X4Neivqp2K4>.

- Riskin, Jessica. *The Restless Clock: A History of the Centuries-Long Argument over What Makes Living Things Tick*. Chicago, IL: University of Chicago Press, 2018. <https://press.uchicago.edu/ucp/books/book/chicago/R/bo21519800.html>.
- Roberge, Jonathan, and Michael Castelle. "Toward an End-to-End Sociology of 21st-Century Machine Learning." In *The Cultural Life of Machine Learning: An Incursion into Critical AI Studies*, edited by Jonathan Roberge and Michael Castelle, 1–29. Cham: Springer International Publishing, 2021. https://doi.org/10.1007/978-3-030-56286-1_1.
- Roberts, Sarah T. "Commercial Content Moderation: Digital Laborers' Dirty Work." *Dirty Work*, 2016, 12.
- Roberts, Sarah T. *Behind the Screen: Content Moderation in the Shadows of Social Media*. Yale University Press, 2019.
- Rosol, Christoph, Benjamin Steininger, Jürgen Renn, and Robert Schlägl. "On the Age of Computation in the Epoch of Humankind." *Nature* 28 (November 2018). <https://www.nature.com/articles/d42473-018-00286-8>.
- Rouvroy, Antoinette, Thomas Berns, and Liz Carey-Libbrecht. "Algorithmic governmentality and prospects of emancipation." *Rezeaux* 177, no. 1 (October 14, 2013): 163–96.
- Rubio, Philip F. "After the Storm: Postal Politics and Labor Relations Following the 1970 U.S. Postal Wildcat Strike, 1970–1981." *Employee Responsibilities and Rights Journal* 30, no. 1 (March 1, 2018): 65–80. <https://doi.org/10.1007/s10672-017-9303-7>.
- Rubio, Philip F. *Undelivered: From the Great Postal Strike of 1970 to the Manufactured Crisis of the U.S. Postal Service*. UNC Press Books, 2020.
- Rumelhart, David E., and James L. McClelland. "Learning Internal Representations by Error Propagation." In *Parallel Distributed Processing: Explorations in the*

- Microstructure of Cognition: Foundations*, 318–62. MIT Press, 1987.
<https://ieeexplore.ieee.org/document/6302929>.
- Santoro, Adam, Andrew Lampinen, Kory Mathewson, Timothy Lillicrap, and David Raposo. “Symbolic Behaviour in Artificial Intelligence.” arXiv, January 21, 2022.
<http://arxiv.org/abs/2102.03406>.
- Saussure, Ferdinand de. *Course in General Linguistics*. Edited by Perry Meisel. Translated by Wade Baskin Edited by Perry Meisel and Haun Saussy. Columbia University Press, 2011.
- Scheele, Carl H. *A Short History of the Mail Service*. Smithsonian Institute Press, 1970.
- Scholz, Trebor. *Digital Labor: The Internet as Playground and Factory*. Routledge, 2012.
- Schultz, Randy. “Compaq to Buy DEC - Jan. 26, 1998.” Accessed January 26, 2020.
<https://money.cnn.com/1998/01/26/technology/compaq/>.
- Seaver, Nick. “What Should an Anthropology of Algorithms Do?” *Cultural Anthropology* 33, no. 3 (August 21, 2018): 375–85. <https://doi.org/10.14506/ca33.3.04>.
- Sejnowski, Terrence J. *The Deep Learning Revolution*. Illustrated edition. Cambridge, Massachusetts: The MIT Press, 2018.
- Sekula, Allan. “The Body and the Archive.” *October* 39 (1986): 3–64.
<https://doi.org/10.2307/778312>.
- Shet, Vinay. “Are You a Robot? Introducing ‘No CAPTCHA ReCAPTCHA.’” Google Online Security Blog (blog), December 3, 2014.
<https://security.googleblog.com/2014/12/are-you-robot-introducing-no-captcha.html>.
- Shet, Vinay. “ReCAPTCHA Just Got Easier (but Only If You’re Human).” Google Online Security Blog (blog), October 25, 2013.
<https://security.googleblog.com/2013/10/recaptcha-just-got-easier-but-only-if.html>.

- Shoemaker, Tyler. "Error Aligned." *Textual Cultures: Texts, Contexts, Interpretation* 12, no. 1 (July 25, 2019): 155–82.
- Silverstein, Michael. "Shifters, linguistic categories, and cultural description." In *Meaning in Anthropology*, edited by Keith H. Basso and Henry A. Selby, 11–55. Albuquerque: University of New Mexico Press, 1976.
- Smith, Brian Cantwell. *On the Origin of Objects*. Cambridge, Mass.: The MIT Press, 1998.
- Smith, Brian Cantwell. *The Promise of Artificial Intelligence: Reckoning and Judgment*. Illustrated edition. Cambridge, MA: The MIT Press, 2019.
- Smolensky, Paul. "On the Proper Treatment of Connectionism." In *Philosophy, Mind, and Cognitive Inquiry: Resources for Understanding Mental Processes*, edited by David J. Cole, James H. Fetzer, and Terry L. Rankin, 145–206. Studies in Cognitive Systems. Dordrecht: Springer Netherlands, 1990.
https://doi.org/10.1007/978-94-009-1882-5_6.
- Snyder, Joel, and Neil Walsh Allen. "Photography, Vision, and Representation." *Critical Inquiry* 2, no. 1 (1975): 143–69.
- Solon, Olivia. "The Rise of 'Pseudo-AI': How Tech Firms Quietly Use Humans to Do Bots' Work." *The Guardian*, July 6, 2018, sec. Technology.
<https://www.theguardian.com/technology/2018/jul/06/artificial-intelligence-ai-humans-bots-tech-companies>.
- Srnicek, Nick. *Platform Capitalism*. Polity Press, 2017.
- Steyerl, Hito. "In Free Fall: A Thought Experiment on Vertical Perspective." *E-Flux* 24. Accessed March 2, 2023. <https://www.e-flux.com/journal/24/67860/in-free-fall-a-thought-experiment-on-vertical-perspective/>.
- Stone, Brad. "Breaking Google Captchas for Some Extra Cash." *Bits Blog* (blog), March 13, 2008. <https://bits.blogs.nytimes.com/2008/03/13/breaking-google-captchas-for-3-a-day/>.

- Suchman, Lucy. "Making Work Visible." *Communications of the ACM* 38, no. 9 (September 1, 1995): 56–64. <https://doi.org/10.1145/223248.223263>.
- Suchman, Lucy. *Human-Machine Reconfigurations: Plans and Situated Actions. Learning in Doing: Social, Cognitive and Computational Perspectives*. Cambridge University Press, 2007.
- Thompson, E. P. *The Poverty of Theory and Other Essays*. New York: Monthly Review Press, 1979.
- Tocqueville, Alexis de. *Democracy in America*. Translated by Harvey C. Mansfield and Delba Winthrop. 1st edition. Chicago, Ill.: University of Chicago Press, 2002.
- Turing, A. M. "Computing Machinery and Intelligence." *Mind*, no. 236 (October 1, 1950): 433–60. <https://doi.org/10.1093/mind/LIX.236.433>.
- United States Postal Service. "Information Security." Information Technology. Accessed March 2, 2023. https://about.usps.com/strategic-planning/cs10/CSPO_12_2010_FINAL_054.htm.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. "Attention Is All You Need." In *Advances in Neural Information Processing Systems*, Vol. 30. Curran Associates, Inc., 2017. <https://papers.nips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- Warner, Julian. "A Labor Theoretic Approach to Information Retrieval." *Journal of the American Society for Information Science and Technology* 59, no. 5 (2008): 731–41. <https://doi.org/10.1002/asi.20782>.
- Weatherby, Leif. "Digital Metaphysics: The Cybernetic Idealism of Warren McCulloch." *The Hedgehog Review*, Springf 2018. <https://hedgehogreview.com/issues/the-human-and-the-digital/articles/digital-metaphysics>.

- Weatherby, Leif. "Intermittent Legitimacy: Hans Blumenberg and Artificial Intelligence." *New German Critique* 49, no. 1 (145) (February 1, 2022): 11–39. <https://doi.org/10.1215/0094033X-9439601>.
- Weinberg, Justin. "Philosophers On GPT-3 (Updated with Replies by GPT-3)." *Daily Nous*, July 30, 2020. <https://dailynous.com/2020/07/30/philosophers-gpt-3/>.
- Wiener, Norbert. *The Human Use Of Human Beings: Cybernetics And Society*. Houghton Mifflin, 1950.
- Wilkes, Emma. "David Guetta Says 'the Future of Music Is in AI' after Eminem Deepfake Vocal Stunt." NME (blog), February 14, 2023. <https://www.nme.com/news/music/david-guetta-says-the-future-of-music-is-in-ai-after-eminem-deepfake-vocal-stunt-3398399>.
- Wilkie, Rob. "Global Networks and the Materiality of Immaterial Labor." In *The Digital Condition*, 50–121. Class and Culture in the Information Network. Fordham University Press, 2011. <https://doi.org/10.2307/j.ctt14brzk3.6>.
- Wimsatt, W. K. *The Verbal Icon: Studies in the Meaning of Poetry*. Revised edition. Lexington: University Press of Kentucky, 1954.
- Winner, Langdon. "Do Artifacts Have Politics?" *Daedalus* 109, no. 1 (1980): 121–36.
- Winograd, Terry. "Understanding Natural Language." *Cognitive Psychology* 3, no. 1 (January 1, 1972): 1–191. [https://doi.org/10.1016/0010-0285\(72\)90002-3](https://doi.org/10.1016/0010-0285(72)90002-3).
- Winston, Brian. "A Mirror for Brunelleschi." *Daedalus* 116, no. 3 (1987): 187–201.
- Wu, Jeffrey, Dario Amodei, Daniella Amodei, Jack Clark, Miles Brundage, Ilya Sutskever, and Alec Radford. "Better Language Models and Their Implications." *OpenAI* (blog). Accessed March 2, 2023. <https://openai.com/research/better-language-models>.
- Wyszomierski, Michael. "Protect Your Site from Spammers with ReCAPTCHA." Official Google Webmaster Central Blog (blog), January 26, 2010.

<https://webmasters.googleblog.com/2010/01/protect-your-site-from-spammers-with.html>.

Zhang, Yang, Haichang Gao, Ge Pei, Shuai Kang, and Xin Zhou. “Effect of Adversarial Examples on the Robustness of CAPTCHA.” In *2018 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery*, 1–109, 2018. <https://doi.org/10.1109/CyberC.2018.00013>.

Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs, 2019.