

Openings and Retrospectives



WHAT SHOULD AN ANTHROPOLOGY OF ALGORITHMS DO?

NICK SEAVER

Tufts University

 <http://orcid.org/0000-0002-3913-1134>

Programmer lore holds that software systems take the shape of the organizations that make them. First formulated as Conway’s Law in the late 1960s, this idea has reappeared periodically since then in various manifestos that reimagine the figure of the programmer and the work of programming. In his *New Hacker’s Dictionary*, Eric S. Raymond (1996, 124) put it this way: “If you have four groups working on a compiler, you’ll get a four-pass compiler.”¹ Social structures emboss themselves onto digital substrates; software is a kind of print left by inky institutions.

Today, technology companies have taken this lesson to heart: if you want to change how your software works, you need to change how the people who make it work. These companies configure their organizational charts in anticipation of software architectures to come. Where old firms stamped out major releases one at a time, the new ones are self-consciously “agile,” rolling out updates continuously and constantly adjusting themselves in response to the data

they now collect in ever greater quantities. Software and the teams that produce it are arranged in small, semiautonomous, readily reconfigurable groups.

These organizational designs often play out in open-plan offices, where desks and employees can be rearranged, aggregated, and disaggregated at will. Facebook, for instance, is largely typed into existence from the world's largest open-office workspace—a massive hangar designed by Frank Gehry that, when it opened, held 2,800 people (Frankel 2015). That enormous room symbolizes at once a corporate commitment to openness, an antihierarchical ethos, and a recursive willingness to rewrite the organizational chart as though it were computer code (cf. Kelty 2005).

Consequently, social architectures, software architectures, and physical architectures echo each other. Walking into a tech company's offices today, one can conjure the feeling of walking into the software itself, as though the building and the people in it—their positions, their interactions, the flow of information—were programs transcoded into social space.

Brad worked at a table by the windows, sitting with the five other members of the radio team. Together, they built playlist-generating algorithms for Whisper, a music recommendation company. Given a song or an album or an artist, Whisper's algorithms could generate endless lists of similar music—what music streaming companies call “radio.” One day, Brad explained to me how he decided if his algorithmic DJ was working.

He *listened* to it, Brad told me. As he tweaked parameters, adjusting the relative weights of features like popularity or tempo, he would try the resulting radio station out, putting in songs and checking the results. If those results sounded good to Brad, he would ask one of his teammates to listen. Then the whole team and their product manager might try it out, adjusting settings along the way. Configurations that seemed plausible and functional would roll out internally to the whole company for feedback, and to small groups of users as tests. If these users seemed to like it, then the configuration would roll out to all of the service's users. At every step, radiating outward from Brad's computer, people are listening. Brad might plug in his favorite artist and decide that the resulting stream is no good. Some unexpected error in the data or software might manifest as a conspicuously out-of-place song. Or, when users unknowingly test the new design, they might listen less, causing the radio team's product manager to roll the change back.

This image has stayed with me as the essence of a contemporary algorithmic system: a steady accumulation of feedback loops, little circuits of interpretation and decision knit together into a vast textile. Every stitch is held together by a moment of human response, a potential rejection shaped by something outside the code, whether it is the arbitrariness of personal preference, the torque of structural bias, or the social force of a formal evaluation framework. Whenever something goes wrong—a public scandal breaks out, a key metric crashes, a manager gets a bad recommendation—the fabric is rewoven, altered by a cascade of human actions. Although this weaving can take many forms, everywhere there are people making little stitches, tacking software to the social world.

ALGORITHMIC DRAMA

Contrast this swatch of sociotechnical fabric with contemporary discourse about algorithms: they have come to “rule our world” (Steiner 2012), not only recommending music but also identifying people as potential criminals, deciding on loan eligibility, trading stocks at unimaginable speed, and shaping the news feeds through which you likely heard about them in the first place. While algorithms are apparently consolidating their power, academic critics find them to be the “key of modern rationality” (Totaro and Ninno 2014), agents of a procedural logic that threatens to displace human judgment, to anticipate and encapsulate us, to draw everyone ever tighter into their surveillant circuits. A pair of popular quips among critical scholars of algorithms captures the general sentiment: if you are not already “slave to the algorithm” (Edwards and Veale 2017), then “the algorithm is gonna get you” (Cheney-Lippold 2017, 93).²

Through the rehearsal of this story, which Malte Ziewitz (2016, 3) calls the “algorithmic drama,” algorithms are figured as powerful, inhuman, and obscure, leaving critics and their readers to fight on behalf of humans. Describing the uncanny effectiveness of recommender algorithms, for instance, Blake Hallinan and Ted Striphas (2016, 126) write: “It is as if the human lifeworld has been invaded by strangers who nonetheless seem eerily adept at navigating foreign land.” Algorithms threaten to blunt culture in their image, to reformat our lives in ways amenable to computation, to strike technical rationality’s killing blow to what remains of the free, serendipitous spirit of human existence (e.g., Gusterson and Besteman, forthcoming).

But what about Brad? He and his coworkers are not strangers to the human lifeworld. Their countless little choices hold algorithms together; they make algorithms responsive to the world; they weave in particular points of view; they

change algorithms over time. While some proponents of algorithms—or machine learning, or artificial intelligence, or whatever complexly responsive software is called by the time you read this—may claim that their systems are autonomous, there are Brads everywhere, tweaking and tuning, repairing and refactoring. In the world of music recommendation, these people often argue that their human sensitivity is, in fact, key to their systems' success: knowing things about music, caring about it, being a culturally conditioned person makes a programmer better at their job, because such things shape their choices and their choices matter. The algorithmic fabric has a human weft.

One might argue that machine learning is different. Now, programmers make software that comes to program itself, learning from data. Maybe we need to consider these algorithms as actants in more-than-human networks, agentic counterparts to their designers, not just their creations; indeed, many people already do this, in more or less social-scientific jargon. But press on any algorithmic decision and you will find many human ones: people like Brad or his manager deciding that a certain error threshold is acceptable, that one data source should be used over another or that a good recommendation means this, not that. These systems are, as a former head of recommendation at Spotify put it, “human all the way down” (Heath 2015). There is no such thing as an algorithmic decision; there are only ways of seeing decisions as algorithmic.

In other words, algorithms are not autonomous technical objects, but complex sociotechnical systems. As I have argued elsewhere, this change in framing reverses many of algorithms' supposedly essential qualities: if the idealized algorithm is supposed to be generic, stable, and rigid, the actual algorithmic system is particular, unstable, and malleable (Seaver 2013). While discourses about algorithms sometimes describe them as “unsupervised,” working without a human in the loop, in practice there are no unsupervised algorithms. If you cannot see a human in the loop, you just need to look for a bigger loop.

Take, for instance, the case of the programmer [Jacky Alciné \(2015\)](#), who discovered that Google Photos had automatically tagged a picture of him and a friend as containing “gorillas.” Alciné and his friend are African American, and their experience is one of the more straightforward and well-known examples of the racist outputs of contemporary algorithmic systems. Alciné and others attributed this algorithmic output to a lack of diversity in the data on which the tagging system was trained and a similar lack of diversity in the programming and testing teams that produced it. While some critics suggested that the algorithm itself could not be said to be racist (in spite of this outcome), this suggestion was a

category error, analogous to the insistence that racism requires intentionally racist individuals (Bonilla-Silva 1997). We can readily see how the worlds of reference of the people working within algorithmic systems—what they anticipate and are sensitive to, their habitus and common sense—come to matter in those systems’ functioning, and thus how diversity in these settings is not only a matter of equitable representation but also one of functional, technical consequence (see Noble 2018).

And we can ask: what happened next? A few hours after Alciné tweeted about the problem, a Google manager named Yonatan Zunger (2015) responded to him, asking for access to Alciné’s data so that the relevant team within Google could investigate the problem. By the next day, the objectionable tag was gone. If we grant algorithms autonomy, treating them as though they are external to human culture, then we cannot make any sense of this story—why it happened in the first place, how it was resolved, and why similar problems might happen again. Tales of autonomous algorithms cannot explain why the system works differently now than it did then. To explain what happened, we need to talk about the makeup of technical teams, the social processes by which those teams define and discover problems, how they identify acceptable solutions, and their culturally situated interpretive processes.³

There is no independent software domain from which computers send forth their agents to insinuate themselves into our cultural worlds. Like any other technology, algorithms can produce unanticipated results or prove challenging to manipulate, but they have no wicked autonomy. Especially in the case of the constantly tweaked and revised algorithms that filter news or make recommendations, these devices work at the pleasure of people who can change them, turn them off, or reorient them toward other goals. These choices emerge from particular cultural worlds, not from some technical outside. Algorithms, we can say, *are* culture—another name for a collection of human choices and practices (Seaver 2017). If we want to understand how they work and especially how they change, we need to attend to people like Brad and the worlds of reference in which they make decisions, not to idealized visions of autonomous algorithmic logics.

This argument is, admittedly, not new: anthropologists have been redescribing technical systems as *sociotechnical* systems for a long time. Bryan Pfaffenberger (1988, 241) states the classic lesson succinctly:

Technology is not an independent nonsocial variable that has an “impact” on society or culture. On the contrary, any technology is a set of social behav-

iors and a system of meanings. To restate the point: when we examine the “impact” of technology on society, we are talking about the impact of one kind of social behavior on another.

Yet it now seems necessary to learn this lesson again, as the rhetorical force of algorithms grows and as they appear to be coming detached from their sociocultural moorings. Arguing for this approach means giving up the excitement of new objects that might demand new methods (and new journals, conferences, and other trappings of academic novelty) in favor of identifying continuities with older concerns.

AGAINST THE ANALOG SLOT

So if the problems posed by algorithms are not new, then what should an anthropology of algorithms do?

When I started asking this question (Seaver 2013), I was struck by the familiar scripts that anthropologists rehearsed when confronted by technical systems that appeared to encroach on our cultural turf. We could offer the analog alternative to algorithms’ digital; the small-scale to their large-scale; the cultural to their technical; the thick description to their thin; the human to their computer. Everywhere, our role in the algorithmic drama seemed already to be cast, such that ethnography was destined to fill whatever negative space algorithms left for us. Whether that made anthropology a complementary counterpart or a perfectly positioned antagonist to algorithmic systems, our role was determined in advance by the very phenomenon we set out to study (Seaver 2015).

I have come to think of this as the anthropology of computing’s version of Michel-Rolph Trouillot’s (1991) “savage slot”—our position within an inherited field of significance that exceeds the discipline and has been integral in shaping it. Let’s call it the *analog slot*: a position nominally opposed to the digital, but which depends on it for coherence. The analog slot is full of the wild and distinctively human remainders that computing can supposedly never consume: our sensitivity to context, our passions and affects, our openness to serendipity and chance, our canny navigation of the cultural field that uncannily algorithms ape but never attain.

The contemporary imprecision of the term *analog* nicely captures the imprecision of the slot, which is effectively a grab bag of qualities thrown together in opposition to computing’s supposed limits. Where *analog* once designated a particular class of computer, it has now come to signify everything outside of computers, a set of objects with no intrinsic coherence. This usage, as Jonathan

Sterne (2016, 32) argues, “comes with a cost, because it effectively diminishes the variety of the world as it elevates conceptions of the digital” (see also Chun 2011). In embracing the analog slot, we do not resist computing so much as grant it definitional supremacy, making it harder to see the variety and detail that might help us to understand how algorithms work and what they do.

Yet a casual analog humanism still dominates anthropological conversations about algorithms and computing more generally. Like other humanisms, this one is selective, carving out a part of the world where it does not apply: the people working within algorithmic systems are, it seems, the only people anthropologists are comfortable treating as *homo economicus*, governed by technical and economic rationality. The kinds of culturally informed decisions I have described so far are thereby rendered invisible.

Rather than endorsing the analog slot, anthropologists should work against it. Trouillot and other critics of anthropological humanisms offer a model. Instead of restating our inherited premises as our conclusions, we can “offer new points of re-entry” and “address directly the thematic field (and thus the larger world) that made (makes) this slot possible” (Trouillot 1991, 40). We do not need to defend the human, as it has been defined in opposition to the algorithmic, nor do we need to build that dichotomy up so that it might be dialectically resolved. Rather, we can refuse it from the start and look for alternatives; we can try to bring that human-versus-algorithm frame into focus and see how it is maintained and reproduced.

Drawing on Marilyn Strathern’s (2005) work on the relations between kinship and knowledge, we can investigate the intellectual kin-making practices that sustain these apparently opposed analog and digital moieties. What ritualized processes of alliance and descent configure anthropology’s relations with its algorithmic alters? How is their connection, stability, and separation maintained? And how do people on the putatively digital side of the divide conceptualize and navigate the discursive field we share with them? What local theories of culture and technology are in play, and where? These are empirical questions that, to date, critics of algorithmic systems have tended to answer with recourse to first principles—to the logic of interacting logics, transcendent of the situations in which ordinary life plays out.

By attending to the ordinary life of algorithmic systems, anthropology can contribute some much-needed empiricism and particularism to critical discourses on algorithms.⁴ Many critiques grant algorithms a homogeneity and coherence that is elusive in practice. We hear that Google search is like predictive policing,

which is like the Netflix recommender, which is like high-frequency trading. These links are supported by claims that certain algorithmic techniques are “domain-independent,” working in any setting. But such claims ignore the practical implementation of algorithms in mundane routines of translation and articulation, maintenance and repair, confusion and ad hoc reasoning. Not all algorithmic systems are the same in their configurations, their effects, or in the worlds of reference of their human minders. There is no more reason to assume that all of these systems share a singular algorithmic logic than there is to assume that all people share a singular human logic.

And while I have focused here on engineers like Brad, they are by no means the only people hidden within supposedly automatic systems. Scholars like [Lilly Irani \(2015\)](#), [Sarah T. Roberts \(2016\)](#), [Hamid Ekbia and Bonnie Nardi \(2017\)](#), and [Mary Gray and Siddharth Suri \(2017\)](#) point our attention to the underpaid and overworked people whose labor powers so-called machine vision or the content moderators who keep obscene content out of our social media feeds. Being human within the machine is not a pleasurable or powerful experience for many of the people thus involved. Again, algorithmic operations are more human than we have been led to believe.

This is one thing an anthropology of algorithms can do: we can reject the casual humanism of the analog slot in favor of finding the people *within* these systems. If Conway’s Law offers a rough emic theory of the relationship between social structure and software, we might fill in the cultural theory: what local practices of analogy, valuation, and contextualization can we find when we recognize that algorithms, in spite of it all, are thoroughly dependent on human sense-making? What cosmologies do people like Brad make their choices in?

Ultimately, anthropology’s greatest contribution to contemporary critical discourse about algorithms may be the corrosive potential of anthropological attention itself, which promises to break down the apparent solidity and coherence of the algorithm. Following [Alfred Gell \(1992, 41\)](#), who originally made this claim about art, we can say that “the ultimate aim of [the anthropology of algorithms] must be the dissolution of [algorithms], in the same way that the dissolution of religion, politics, economics, kinship, and all other forms under which human experience is presented to the socialized mind, must be the ultimate aim of anthropology in general.”

NOTES

Acknowledgments Thanks to Christina Agapakis, Adia Benton, Ian Lowrie, the *Cultural Anthropology* editorial staff, and the anonymous reviewers for their suggestions about what I should do with this piece.

1. A compiler is a program that converts human-readable source code into machine code that can be executed by a computer. A four-pass compiler would perform this conversion in four sequential passes; Raymond is suggesting that having four groups working on this problem would result in the problem being divided into four parts.
2. These lines play on the titles of two 1980s pop music hits: Grace Jones's "Slave to the Rhythm" (1985) and Gloria Estefan's "Rhythm is Gonna Get You" (1987).
3. What happened, according to a recent report in *Wired* (Simonite 2018), was that Google engineers simply blocked the system from labeling most nonhuman primates. This particular solution, blunt as it is, resulted from a human decision-making process and it meaningfully changes how "the algorithm" works, even if it seems too crude to count as machine learning.
4. Methodological challenges face the would-be ethnographer of algorithmic systems, but again, these are only apparently new: anthropologists have tried to study the cultural worlds of elite and secretive groups for a long time, and we already have many useful tactics for trying to understand these worlds when access is elusive or complex (e.g., Gusterson 1997; Ortner 2010; Jones 2011; Mahmud 2014; Seaver 2017).

REFERENCES

- Alciné, Jacky (@jackyalcine)
 2015 "Google Photos, y'all fucked up. My friend's not a gorilla." Twitter post, June 28. <https://twitter.com/jackyalcine/status/615329515909156865>.
- Bonilla-Silva, Eduardo
 1997 "Rethinking Racism: Toward a Structural Interpretation." *American Sociological Review* 62, no. 3: 465–80. <https://doi.org/10.2307/2657316>.
- Cheney-Lippold, John
 2017 *We Are Data*. New York: New York University Press.
- Chun, Wendy Hui Kyong
 2011 *Programmed Visions: Software and Memory*. Cambridge, Mass.: MIT Press.
- Edwards, Lilian, and Michael Veale
 2017 "Slave to the Algorithm? Why a 'Right to an Explanation' Is Probably Not the Remedy You Are Looking For." *Duke Law and Technology Review* 16. <https://scholarship.law.duke.edu/dltr/vol16/iss1>.
- Ekbia, Hamid R., and Bonnie A. Nardi
 2017 *Heteromation, and Other Stories of Computing and Capitalism*. Cambridge, Mass.: MIT Press.
- Frankel, Todd C.
 2015 "What These Photos of Facebook's New Headquarters Say about the Future of Work." *Washington Post*, November 30. <https://www.washingtonpost.com/news/the-switch/wp/2015/11/30/what-these-photos-of-facebooks-new-headquarters-say-about-the-future-of-work>.
- Gell, Alfred
 1992 "The Technology of Enchantment and the Enchantment of Technology." In *Anthropology, Art, and Aesthetics*, edited by Jeremy Coote and Anthony Shelton, 40–63. New York: Oxford University Press.
- Gray, Mary L., and Siddharth Suri
 2017 "The Humans Working Behind the AI Curtain." *Harvard Business Review*, January 9. <https://hbr.org/2017/01/the-humans-working-behind-the-ai-curtain>.

- Gusterson, Hugh
 1997 "Studying Up Revisited." *Political and Legal Anthropology Review* 20, no. 1: 114–19. <https://doi.org/10.1525/pol.1997.20.1.114>.
- Gusterson, Hugh, and Catherine Besteman, eds.
 Forthcoming *Robohumans: How Algorithms are Remaking Social Life*. Chicago: University of Chicago Press.
- Hallinan, Blake, and Ted Striphas
 2016 "Recommended for You: The Netflix Prize and the Production of Algorithmic Culture." *New Media and Society* 18, no. 1: 117–37. <https://doi.org/10.1177/1461444814538646>.
- Heath, Alex
 2015 "Spotify is Getting Unbelievably Good at Picking Music—Here's an Inside Look at How." *Business Insider*, September 3. <http://www.businessinsider.com/inside-spotify-and-the-future-of-music-streaming>.
- Irani, Lilly
 2015 "Justice for 'Data Janitors.'" *Public Books*, January 15. <http://www.publicbooks.org/justice-for-data-janitors>.
- Jones, Graham M.
 2011 *Trade of the Tricks: Inside the Magician's Craft*. Chicago: University of Chicago Press.
- Kelty, Christopher
 2005 "Geeks, Social Imaginaries, and Recursive Publics." *Cultural Anthropology* 20, no. 2: 185–214. <https://doi.org/10.1525/can.2005.20.2.185>.
- Mahmud, Lilith
 2014 *The Brotherhood of Freemason Sisters: Gender, Secrecy, and Fraternity in Italian Masonic Lodges*. Chicago: University of Chicago Press.
- Noble, Safiya Umoja
 2018 *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press.
- Ortner, Sherry B.
 2010 "Access: Reflections on Studying up in Hollywood." *Ethnography* 11, no. 2: 211–33. <https://doi.org/10.1177/1466138110362006>.
- Pfaffenberger, Bryan
 1988 "Fetishized Objects and Humanized Nature: Toward an Anthropology of Technology." *Man N.S.* 23, no. 2: 236–52. <https://doi.org/10.2307/2802804>.
- Raymond, Eric S.
 1996 *The New Hacker's Dictionary*. 3rd edition. Cambridge, Mass.: MIT Press.
- Roberts, Sarah T.
 2016 "Commercial Content Moderation: Digital Laborers' Dirty Work." In *The Intersectional Internet: Race, Sex, Class and Culture Online*, edited by Safiya Umoja Noble and Brendesha M. Tynes, 147–59. New York: Peter Lang.
- Seaver, Nick
 2013 "Knowing Algorithms." Paper presented at Media in Transition 8, Cambridge, Mass., May 4.
 2015 "Bastard Algebra." In *Data! Now Bigger and Better*, edited by Tom Boellstorff and Bill Maurer, 27–46. Chicago: Prickly Paradigm Press.
 2017 "Algorithms as Culture: Some Tactics for the Ethnography of Algorithmic Systems." *Big Data and Society* 4, no. 2. <https://doi.org/10.1177/2053951717738104>.
- Simonite, Tom
 2018 "When It Comes to Gorillas, Google Photos Remains Blind." *Wired*, January 11. <https://www.wired.com/story/when-it-comes-to-gorillas-google-photos-remains-blind>.
- Steiner, Christopher
 2012 *Automate This: How Algorithms Came to Rule Our World*. New York: Portfolio.

- Sterne, Jonathan
 2016 "Analog." In *Digital Keywords: A Vocabulary of Information Society and Culture*, edited by Benjamin Peters, 31–44. Princeton, N.J.: Princeton University Press.
- Strathern, Marilyn
 2005 *Kinship, Law, and the Unexpected: Relatives are Always a Surprise*. New York: Cambridge University Press.
- Totaro, Paolo, and Domenico Ninno
 2014 "The Concept of Algorithm as an Interpretative Key of Modern Rationality." *Theory, Culture, and Society* 31, no. 4: 29–49. <https://doi.org/10.1177/0263276413510051>.
- Trouillot, Michel-Rolph
 1991 "Anthropology and the Savage Slot: The Poetics and Politics of Otherness." In *Recapturing Anthropology: Working in the Present*, edited by Richard G. Fox, 17–44. Santa Fe, N.M.: School of American Research Press.
- Ziewitz, Malte
 2016 "Governing Algorithms: Myth, Mess, and Methods." *Science, Technology, and Human Values* 41, no. 1: 3–16. <https://doi.org/10.1177/0162243915608948>.
- Zunger, Yonatan (@yonatanzunger)
 2015 "@jackyalcine Thank you for telling us so quickly! Sheesh. High on my list of bugs you *never* want to see happen. ::shudder::" Twitter post, June 29. <https://twitter.com/yonatanzunger/status/615383011819794432>.