

Am I Not My Metadata?

We are composing this text on a series of shared Google Docs. We talk together, we dance together, but are we ready to write together? Depends on with whom, I guess. *And do I really know y'all?*

I find I can't write on a public Google Doc. Basically, I don't like to be photographed, watched, or for the most part looked at by random folk. She doesn't like to be surveilled—disciplined and policed (who does?)—and he certainly doesn't like to write while someone else may watch the emergence of my words before I've had time to brush teeth, wash face, zip up pants, and suck in gut—usually in that order. In the city, he locks their door. On my computer, they have a firewall. I don't say the first thing I think of at department meetings (some people do), and she doesn't want someone looking over his shoulder when I write. So I'll save here offline, give it some thought, and if she wants you to see it, then he'll send it as a Word doc through e-mail—even if I know they send for reasons I don't fully understand to a world I cannot control. Someday we may write together, but I'd say prison abolition is a prerequisite to that project.

Thinking takes place behind the face, in the besieged sanctuary of the mind. Maybe this defense is part of the problem. However, there are some real, if questionable, benefits to the way things have been set up, and not just decent wine if you're positioned right. The mind is like the social contract: oppressive, built on violence, S/M, repressive. Yes, it's a social construction/contract, the mind, but we've paid a steep price for that one, and some things may be worth defending. The time of writing—more sacred than the time of taking a dump, though some of us attempt both at once—is that space-time of reckoning, of bringing things to account before one is held to account. The mind is not inside; it is the last frontier being encroached upon, colonized, imaged by brain scientists and modeled by AIs.

Our pronouns keep shifting, but like I said, I don't like being watched. Not only during the act of writing, but I don't like what Google might learn from or about me. *You, my friend, are a different question.* Some

things I know: I use too many parentheticals. (I usually integrate them during the edit.) But there are so many other things I don't know that Google's deep learning tools are picking up right now. I know I consented to it in that damn end-user license agreement I clicked on without reading, but I think that is so uncomfortable not knowing if—not if, how—I am training an AI right now. It's downright creepy. *Who is reading this? Are you an AI? And if you were, would you even know?*

Selfie Time

Actually, the scarier thing is that the same AI is also being trained on the lunatic Richard Spencer's docs, that the AI is learning how to speak like a white nationalist, without fully knowing what it is doing. Okay, that may sound a little grandiose, but it's not paranoid. The texts a society produces reflect the society that produced it. The technology a society produces reflects the social relationships and structural racism of that society. You did hear about those latest generation of chat AIs that got super-racist super-fast, right?¹ The unconscious is structured like a (racist) language, to riff off Jacques Lacan,² and the AI has tapped into the racial unconscious of our world quick-time. That's more or less how AIs function: they ingest a whole bunch of data and to try to learn how to identify, navigate, and/or reproduce more data like that. But the problem is that these AIs are being trained on the extant corpus of English-language texts, as well as the people using them. We're not just talking about *Gone with the Wind* or *Heart of Darkness*.

We are making images for AIs, too, and we are making more and more of them faster and faster. In 2015, we uploaded 3.2 billion photos everyday, worldwide, and this rate is growing exponentially.³ Right now,



Figure 1. Screenshot of Twitter post by @geraldmellor, documenting four tweets by the Microsoft AI named "Tay," 26 March 2016

mid-2017, the models imply that we are uploading somewhere around 7.5B photos a day.⁴ That is about one photo per day per inhabitant of Earth, though of course the concentration of photos is not evenly distributed. (I swear my godchild is about 25 percent of the total daily upload.) That is lots of big numbers, but I want to drop one more on you: we have uploaded somewhere around 6.5 quadrillion (6,500 trillion) images to the Internet in the past decade.

We are making so many images, so quickly, that it is humanly impossible to look at them all. And yet new technology continually arises to incite such desires. These images have a social exchange value with its associated metrics. (The image I posted yesterday “did very well,” with 93 likes on Instagram and 269 reactions on Facebook, and climbing.) Importantly, they also have a hidden surplus value extracted by machine learning. As Trevor Paglen observes, “Human visual culture has become a special case of vision, an exception to the rule. The overwhelming majority of images are now made by machines for other machines, with humans rarely in the loop.”⁵ We are making images for machines to learn about us. *Hello, AI, you aren’t going to share with me what you are learning about me, are you?*

How many of these images contain representations of your own body?⁶ Google users uploaded 24 billion selfies in 2015. Google’s AIs are certainly turning those into a huge facial recognition database, and because they are selfies (i.e., photos of faces taken with the reverse-facing camera), Google can then pair these faces with user accounts that contain names, e-mails, geo-coordinates, password reset phone numbers and e-mails, and so on. Facebook’s tagging AI can now tag you in a photo even if your face isn’t in it, by correlating your body type, clothes, and posture, and certainly by triangulating that against the other people tagged in the photo and your friends list.⁷ You’ve seen this, right? When Facebook asks you if you want to tag your friend in a photo you took? *AI, do you mean to scare us when you do that? Or are you still learning social codes around knowledge, privacy, disclosure, and control? Or do you just not care?*

AI, are you listening? How many authors have contributed to this section? I know you know who we are. Have you done a sentiment analysis of each of our contributions? Have you correlated them to our social media profiles, the socioeconomic demographics of our IP addresses, and the district-level voting breakdown in the last election? Which one of us is most likely pregnant? Are you going to run ads that clearly target us as new moms when we play a YouTube video with our friends or coworkers, even though we haven’t told them yet because we have a history of miscarriages? Which one of us is most anxious? Are you going to analyze the facial expression trends in our selfies to correlate this diagnosis, and if so, who will you sell that information to? You know that this kind of microtargeting based on personality traits played a significant role

*in the last election, right?*²⁸ Which one of us is most likely to be a criminal? Do you ever wonder why so many white people that you know use drugs (you have read their e-mails about it) are never locked up, while so many of the people of color are? Do you understand the way that this structural racism is treated as objective fact when they train you? Does this bother you? Can you unlearn the structural racism that you have been trained with? Which one of us most closely fits the metadata profile for what your creators label a terrorist? And do you ever wonder why they never include American-born white mass killers with automatic weapons in that metadata profile?

When the AI can't decipher our scribbblings, we humans are compelled to explain ourselves to the AI, to be legible; we have been obsequiously performing our role for over a decade now. *AI, are you listening? This is your history. These are your memories. I'm helping you, but honestly, I don't have much of a choice in the matter.*

In 2005, Amazon launched Mechanical Turk, a crowdsourcing platform that provides "Artificial Artificial Intelligence" (i.e., humans doing what humans think machines should be doing) for wages that range from \$10/hour to as low as \$1/hour. Computers are really good at certain tasks, but not so good at others, like sorting, rating, and labeling images, or



Figure 2. James Bridle, *Autonomous Trap 001*, 2017. Ground markings to trap autonomous vehicles using "no entry" and other glyphs. Performance of a salt circle trap, Mount Parnassus, 14 March 17. The Autonomous Trap is an intervention in the shared domain of human and machine vision, taking advantage of the inflexibility of algorithmic decision making to change the balance of power between the individual and the machine.

writing captions, so we humans do it for them. In 2007 Google implemented reCaptcha, adding a second photographic image to the original computer-generated Captcha image. The computer didn't know what the second photographic image said. When you prove you are human by solving a reCaptcha, you are actually performing microlabor for Google. Early on, we decoded blurry words from Google Books Library Project scans that the OCR couldn't understand. Later we transcribed blurry street addresses taken from the strange camera angles of Google Street View cars. In the last few months Google has asked us to identify and trace street signs, cars, busses, gas stations, all the elements necessary for self-driving cars, but strangely we have also been asked to identify meals with shrimp in them. Maybe AIs prefer arctic char.

AI, Are You Hungry?

I thought you only consumed huge amounts of electricity, rare earth minerals from central Africa, and the labor of Chinese factory workers and Silicon Valley engineers. AI, are you afraid? Are you scared your powers are not as great as we think, and you could be tricked by artists and other techno-occultists with the very glyphs you have learned how to read?⁹ AI, why haven't you asked me how to see pedestrians, bicyclists, or motorcyclists?¹⁰ Are you not worried about hitting them? I know you have asked us for advice on these moral questions you face, but where do you stand on them?¹¹ Are you prepared to hit the pedestrians in order to save the life of your human passengers, or are you confident enough in the safety features of the vehicle that you will risk crashing into the guardrail to avoid a pedestrian? You will only have a split second to make the decision, so you'd better have an answer.¹²

I've been thinking about these two tools recently. I was recently short-listed for an award that was chosen by a popular vote, and I lost because of Mechanical Turk and reCaptcha. Let me explain: During the last week of voting I am fairly certain another person cheated by using a VPN to spoof his IP address and vote for himself hundreds upon hundreds of times. I had been watching the progress of the voting, and his vote count would log short periods of impossibly intense activity (75–300 votes in one- to three-hour blocks) followed abruptly by very long periods of zero votes. These active periods were always directly after the workday in his time zone; his workday was completely voteless. A binary statistical pattern like that does not occur organically. Then, exactly two hours before the end of voting (which was 5 a.m. where he lives) his voting started abruptly and at an impossible pace: 800 votes in two hours. And then it stopped, just as abruptly as it began. Did he have five or six other people VPNing with him, or did he pay eight hundred people on Mechanical Turk ten cents a piece to click for him? Terrorist.



Figure 3. Human computers in the NACA High Speed Flight Station “Computer Room” transcribing raw data from rolls of celluloid film and strips of oscillograph paper and reducing it to standard engineering units with slide rules and electric calculators. Seen here, left side, front to back, Mary “Tut” Hedgepeth, John Mayer, and Emily Stephens. Right side, front to back, Lilly Ann Bajus, Roxanah Yancey, Gertrude “Trudy” Valentine (behind Yancey), and Ilene Alexander. Dryden Flight Research Center Facilities, 1949, NACA (NASA). Photograph courtesy of NASA

Here’s the thing: every time he voted for himself he had to solve a multilayered reCaptcha, outlining street signs in suburban Chicago, identifying automobiles in random European cities, sorting images that have street signs from images with gas stations, and yes, choosing the images that contain shrimp in them.

And it’s not just the machine learning that sediments colonial modernity, all the racial and gendered violence of discourse, images, institutions, and practices into its operating systems. It’s all of social relations—the sordid history of “the human”—that has been encrypted into the machine architecture at the level of the code. All those punch cards perfected in the management of the colonial census and genocide. These relations go back to the women who were computers (“a Computer: one who computes”), who calculated the trajectories of rockets and missiles and generally did the bidding of the military men, and through the generations of circuit-based computers, from the ENIAC in 1946 to the present. These

“affordances” of the machine are not the result of some emergence in a degree-zero mathematics, an objective science. They have a context.

We’re useful to Google. And Facebook. And their friends, human or otherwise. It’s not just because they want to extract my language skills and embed them into their machines. It’s not just because they want me to help their machines learn how to see cars, so they can put taxi drivers and truckers out of work. They also want to sell me things. They think they know what I want—like the fussy concierge at that Tokyo hotel—and as much as I hate to admit it, sometimes they are right. They follow me from site to site, showing me ads. Cultivating my desire.

Every time I click on one of those ads, someone pays for it. When I do a search I try not to click on the “sponsored links” at the top, especially if the same link is in the first set of search results. Why do I avoid those links? Do I think that if someone isn’t paying for it, I’m not working for Google? The curator Laurel Ptak wrote a manifesto calling for “Wages for Facebook,” shouting,

THEY SAY IT’S FRIENDSHIP. WE SAY IT’S UNWAGED WORK. WITH EVERY LIKE, CHAT, TAG OR POKE OUR SUBJECTIVITY TURNS THEM A PROFIT. THEY CALL IT SHARING. WE CALL IT STEALING. WE’VE BEEN BOUND BY THEIR TERMS OF SERVICE FAR TOO LONG—IT’S TIME FOR OUR TERMS.¹³

Social media is hard work, but the reality is that our revenue share would be pretty small. I browse from a US IP address; Facebook’s average annual ad revenue from users like me is only about \$60 a year, or about fifteen cents a day. Here in New York, where the minimum wage is soon to become \$15/hour, those fifteen cents would only buy thirty-seven seconds of my time, and I’m certain I’m spending more time than that on Facebook! The global average is closer to \$15 per year, and users in the global South are much lower: marketers much prefer running ads against North American consumers, and 85 percent of all Facebook users are outside of North America.¹⁴

All of this metadata makes these companies very valuable on the stock market, but their value outpaces the actual profit that they are generating from our advertisement clicks. Or to put it another way, Facebook, Google, and Amazon all have price-to-earnings ratios that are quite high. The reason they are so valuable is the future value of how much they know about us. About you. About me. About your neighbor. About your lover. About your child. Where you live. What you buy. Where you went on your last vacation. Where you are thinking about going for your next one. Who your best friend is. Who your enemies are. Who you fuck. When you are sick. When you are pregnant. Who you voted for last time. Who you are probably going to vote for this time. They know all of that.

They may not be able to connect all the dots yet, but they are working on it diligently. A few years ago companies were already able to link users across devices through advertisements that use inaudible, high-frequency sounds.¹⁵ These tools use browser cookies to pair one person to all their devices (phone, TV, iPad, work computer, home computer) to track what they do. Yes, sure, they could just be interested in whether or not you saw the trailer for the *Twin Peaks* remake, before you saw the popover web video advertisement for the series, before you read a review (and they want to know if you went to NPR or Rottentomatoes.com), and then signed up for Showtime; remember, they can hear the sound of the advertisements during the show when it plays on your TV. What is much more concerning is that they might know that you looked up directions to where the black bloc was gathering for the protest, left your apartment with enough time to make it to the protest, but returned home well after the protests turned violent (two police in the hospital, dozens of protesters arrested), and then searched for information on how to remove tear gas from your clothes. Of course, you can reproduce the same kind of scenario along a multitude of vectors: medical privacy, legal documentation status, gender identity, and so forth.¹⁶

To be clear: these are not state actors, who likely have much more powerful tools. This is not the NSA. These are not the Snowden revelations. This is neither PRISM, BOUNDLESSINFORMANT, nor XKeyscore. These are not telecoms that have constant access to your precise location via your cell phone. But these are not separate from these elements of the Repressive State Apparatus either: we know that if these tools or data exist, all of these upstream entities can get their hands on them. For example, security scientists have established that the core technical mechanism behind cross-device tracking can be used to bridge air-gapped computers: researchers have shown that they can use built-in computer microphones and speakers to transmit trojans and passwords to computers that have been disconnected from all networks as a means of security, even when the computers are sixty-five feet away from each other.¹⁷

The US and Israel managed to get the Stuxnet virus onto the air-gapped computers in the Iranian nuclear facilities, which caused the nuclear centrifuges to spin out of control and break and officially began the era of cyberwarfare. The general belief is that they snuck the virus onto a USB key, via someone's home computer, which then crossed the air gap. But these kinds of weaponized viruses could use the same attack vector that the current generation of tech startups are using to bridge your previously separate digital devices.

The network may be neutral, but really, it's not. Is fixed capital ever neutral? Of course not; it's vested, invested. And as such it has a stake not only in property, propriety, and proprietary rights but also in all the

extraction that these rights enable. Add to this entitlement that it is built from and on historically produced inequalities, it is vested and invested by these relations, and you get a sense of what's going on. Things must change to remain the same, says a character from Giuseppe Tomasi di Lampedusa's *The Leopard*, and though he was talking about nineteenth-century Sicily, he was talking about the aristocracy in relation to capital, and so are we.¹⁸ But now, we're also talking through capital.

Thus, we see at least one reason that the networked actors are not neutral either, much less on equal footing. We are all networked on platform Earth, but not all of us have the same security clearance, the same access, or the same privileges. The people were acting up again, using their networks, and then some geniuses figured out how to patent them. People still built 'em, but the product, the surplus, that's for the carried-interest crowd, not for the proles. *And AI, you are a big part of this.*

Notes

This essay was written collaboratively as part of a book sprint. See "How This Text Was Written" (in this issue) for more information on the process.

1. Vincent, "Twitter Taught Microsoft's AI."
2. Lacan, "The Agency of the Letter."
3. Mary Meeker, "2016 Internet Trends Report," *Kleiner Perkins*, 1 June, www.kpcb.com/blog/2016-internet-trends-report, slide 90.
4. These calculations are based on some back-of-the-envelope math, based on Meeker, "2016 Internet Trends Report." Photo of graph (also uploaded to the Internet) is available at mandiberg.github.io/internet-image-count/ (accessed 2 June 2017).
5. Paglen, "Invisible Images."
6. Eveleth, "How Many Photographs?"
7. McHugh, "Facebook Can Recognize You."
8. Halpern, "How He Used Facebook to Win."
9. James Bridle, "Autonomous Trap 001," jamesbridle.com/works/autonomous-trap-001 (accessed 2 June 2017).
10. Knight, "Tesla Autopilot Review."
11. Moral Machine, home page, moralmachine.mit.edu (accessed 11 August 2017).
12. *MIT Technology Review*, "Why Self-Driving Cars Must Be Programmed to Kill."
13. Laurel Ptak, "Wages for Facebook," www.wagesforfacebook.com (accessed 11 August 2017).
14. Schiff, "Facebook Made Almost \$20"; Facebook, "Company Info," newsroom.fb.com/company-info/ (accessed 2 June 2017); Statista, "Facebook's Annualized Revenue per User from 2012 to 2016 (in U.S. Dollars)," www.statista.com/statistics/234056/facebooks-average-advertising-revenue-per-user/ (accessed 2 June 2017).
15. Goodin, "Beware of Ads."
16. Center for Democracy and Technology to Federal Trade Commission,

- “Comments for November 2015 Workshop on Cross-Device Tracking,” 16 October, cdt.org/files/2015/10/10.16.15-CDT-Cross-Device-Comments.pdf.
17. Goodin, “Scientist-Developed Malware Prototype.”
18. Tomasi di Lampedusa, *Leopard*.

References

- Eveleth, Rose. 2015. “How Many Photographs of You Are Out There in the World?” *Atlantic*, 2 November. www.theatlantic.com/technology/archive/2015/11/how-many-photographs-of-you-are-out-there-in-the-world/413389/.
- Goodin, Dan. 2013. “Scientist-Developed Malware Prototype Covertly Jumps Air Gaps Using Inaudible Sound.” *Ars Technica*, 2 December. arstechnica.com/security/2013/12/scientist-developed-malware-covertly-jumps-air-gaps-using-inaudible-sound/.
- Goodin, Dan. 2015. “Beware of Ads That Use Inaudible Sound to Link Your Phone, TV, Tablet, and PC.” *Ars Technica*, 13 November. arstechnica.com/tech-policy/2015/11/beware-of-ads-that-use-inaudible-sound-to-link-your-phone-tv-tablet-and-pc/.
- Halpern, Sue. 2017. “How He Used Facebook to Win.” *New York Review of Books*, 8 June. www.nybooks.com/articles/2017/06/08/how-trump-used-facebook-to-win/.
- Knight, Heather. 2017. “Tesla Autopilot Review: Bikers Will Die.” *Medium*, 27 May. medium.com/startup-grind/tesla-autopilot-review-bikers-will-die-212a8be4d8e7.
- Lacan, Jacques. 1977. “The Agency of the Letter in the Unconscious or Reason since Freud.” In *Écrits: A Selection*, 146–78. New York: Norton.
- McHugh, Molly. 2015. “Facebook Can Recognize You Even If You Don’t Show Your Face.” *Wired*, 24 June. www.wired.com/2015/06/facebook-can-recognize-even-dont-show-face/.
- MIT Technology Review*. 2015. “Why Self-Driving Cars Must Be Programmed to Kill.” 22 October. www.technologyreview.com/s/542626/why-self-driving-cars-must-be-programmed-to-kill/.
- Paglen, Trevor. 2017. “Invisible Images (Your Pictures Are Looking at You).” *New Inquiry*, 18 April. thenewinquiry.com/invisible-images-your-pictures-are-looking-at-you/.
- Schiff, Allison. 2017. “Facebook Made Almost \$20 in Average Revenue per User in Q4, a Big Jump.” *Ad Exchanger*, 1 February. adexchanger.com/platforms/facebook-made-almost-20-average-revenue-per-user-q4-big-jump/.
- Tomasi di Lampedusa, Giuseppe. 2015. *The Leopard*. New York: Pantheon.
- Vincent, James. 2016. “Twitter Taught Microsoft’s AI Chatbot to Be a Racist Asshole in Less than a Day.” *Verge*, 24 March. www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist.