

From Enclosure to Foreclosure and Beyond: Opening AI's Totalizing Logic

1. Introduction

Over the past few years, numerous researchers have pointed to the violently extractivist logic of current AI technology, both in terms of the material resources needed for the hardware and the labor that goes into data gathering and labeling (Crawford and Joler 2018; Couldry and Mejias 2019a, b; Verdegem 2022; Morreale et al. 2023). This logic has been described as a new form of colonialism. Nick Couldry and Ulises Mejias describe “data colonialism” in terms of data relations that normalize “the exploitation of human beings through data, just as historic colonialism appropriated territory and resources and ruled subjects for profit” (Couldry and Mejias 2019a, p. 336).¹ Encompassing not only data but machine learning models as such, Louise Amoore describes their assemblage as “colonizing in ways that incorporate ever-increasing layers, [and that] extend to ever more domains of life . . .” (Amoore 2022, p. 15). The awareness for the extractive, appropriative, and exploitive character of digital technology, however, is not new. In 2000 while the Internet was still in its nascent stage, Tiziana Terranova analyzed the ways in which capital appropriates free digital labor to extract value (Terranova 2000). She argues that unlike what a more orthodox Marxist interpretation of the Internet would lead us to believe, capitalism in the digital economy does not proceed by appropriating and managing digital labor directly. Instead, it encourages and promotes free digital labor—free because freely carried out *and* provided for free—to then appropriate—or enclose—the value it produces. Terranova’s analysis, while addressing a much earlier stage of the digital economy, continues to be highly relevant for understanding current AI: Large language models like GPT are trained on a vast amount of data produced through free labor—among others on Wikipedia and Reddit—to then make the tool available for a price. More recently, Shoshana Zuboff has shown how surveillance capitalism based on digital technologies repeats the “‘original sins’ of primitive accumulation” by extracting and capitalizing on the “behavioral surplus” that we produce in our day-to-day interaction with digital technologies (Zuboff 2019, p. 11).

In the present paper, I want to reframe the issue of appropriation, extraction, and dispossession in terms of the operations of enclosure and foreclosure. Doing so allows me to describe the internal logic of current AI—an assemblage of machine learning models trained on big data. While enclosures are the product of a well-studied set of operations pertaining to both the constitution of the sovereign State and the primitive accumulation of capital, which comprises colonization, here, I want to recover an older form of the enclosure operation to then contrast it with foreclosure in order to gain a better understanding of the effects of current algorithmic rationality. I argue that the act of enclosing is to be understood as a set of fundamental operations that consist in producing structural distinctions between

¹ The authors insist that “data colonialism” is not to be understood as a metaphor but as “a highly distinctive exercise of power” (Couldry and Mejias 2019b, p. xi).

1
2
3
4
5
6
7
8
9
10
11
12
inside and outside, inclusion and exclusion—whether by drawing lines on a map, constructing border walls, or by algorithmically categorizing and (mis)recognizing people and things. Tracking the transformation of an enclosure-logic into one of foreclosure, I show how current AI perpetuates and at the same time expands forms of extraction and dispossession toward a totalizing horizon. For, while the outside is essential and constitutive to the logic of enclosure,² foreclosure, by contrast, is characterized by the “totalizing desire” (Ring 2014) to not leave anything out of it. At the same time, I contend that AI-technology can never live up to its foreclosing propensity, as it is necessarily materially and logically limited.

13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
After providing a short genealogy of the concepts of enclosure and foreclosure, I analyze what I contend are the three essential components of current AI—the dataset, the machine learning model, and the subject—from the perspective of the shift from a logic of enclosure to one of foreclosure. I first compare big data’s architecture to its older correlative of the archive; I then interrogate the changing status of the outlier from statistical modeling to machine learning models; finally, I discuss the generative large language model ChatGPT’s claim to “not have a position” while speaking in the first person singular pronoun. I show through this analysis that while enclosure is a set of perceptible and explicit operations, making enclosures susceptible to contestation, AI’s logic of foreclosure expands the enclosing operations of selection, exclusion, and dispossession toward a totalizing horizon while at the same time invisibilizing them, thus making AI technologies harder to resist and act against.

31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 64 65 2. Enclosure and Foreclosure: A Genealogy

2.1 Enclosure

The concept of *enclosure* is usually addressed by two distinct yet interacting fields of research: political theory and economics. In political theory, the enclosure designates both the origin and the condition of State sovereignty. Taking on Carl Schmitt, Wendy Brown describes sovereignty as the result of the operation of enclosure, that is, the spatial division of the Earth—by marking the borders of a territory with milestones, on a map, or by means of border walls: “. . . it is through the walling off of space from the common that sovereignty is born” (Brown 2010, p. 57). As Brown insists, to enclose is both temporally and logically conceived as the condition for sovereignty in classical political theories (Brown 2010, p. 59). Focusing on today’s nation-states, she contends that enclosures appearing in the form of border walls function as a theatrical affirmation of sovereignty at the very moment when nation-states’ power is waning.

In Marxist and post-Marxist economic theories, the act of *enclosing* is viewed as a set of operations pertaining to primitive and secondary accumulation. Primitive accumulation consists in the enclosure or appropriation of land held in common or in the colonizing of foreign land (Marx 1976). Secondary

² The demonstration of the necessary character of the constitutive other to the production and maintenance of the autonomous and sovereign self is one of the central theoretical gains of Feminist, Postcolonial and Black Studies.

1 accumulation results from the production of an exteriority to capital, which can then be reappropriated
2 or re-enclosed to extract value through dispossession. To name but one example, what Marx designates
3 as “surplus population,” that is, for instance, a population pushed out of the labor force through
4 technological advances, enables the process of secondary accumulation or what David Harvey calls in
5 his book *The New Imperialism* “accumulation by dispossession” (Harvey 2003). The production of a
6 surplus, jobless population rejected outside of capital allows capital to put workers under pressure to
7 reduce labor costs and thus increase surplus-value. As a constitutive operation of capital, accumulation
8 by dispossession is not an outdated form of accumulation but, as Sociologist Klaus Dörre compellingly
9 shows, an ongoing one (Dörre 2015).

10 While existing discussion on enclosure generally focuses on either one of the two perspectives, I
11 contend that they share a common ancestry in the form of a cultural technique. A cultural technique
12 designates a chain of operations producing distinctions fundamental to how a culture understands itself
13 (Siegert 2013, p. 61). To enclose is a fundamental cultural technique, as Cornelia Vismann contends in
14 her paper *Cultural Techniques and Sovereignty* (2013). Enclosures, in both a political and economic
15 sense, emerge from—are made possible in the first place by—marking the ground with a plough. She
16 writes:

17
18
19
20
21
22
23
24
25
26
27
28 To start with an elementary and archaic cultural technique, a plough drawing a line in the ground: the agricultural
29 tool determines the political act; and the operation itself produces the subject, who will then claim mastery over
30 both the tool and the action associated with it. Thus, the *Imperium Romanum* is the result of drawing a line – a
31 gesture which, not accidentally, was held sacred in Roman law. Someone advances to the position of legal owner
32 in a similar fashion, by drawing a line, marking one’s territory – ownership does not exist prior to that act
33 (Vismann 2013, p. 84).³

34
35 For Vismann, the subject is only *de jure* sovereign. It is media, understood as the product of a cultural
36 technique and the “operational script” inscribed in it, which determines the field of possible actions for
37 a subject. It is only in retrospect that the operation of drawing a line is enshrined in the law as “sacred.”
38 Here, I wish to understand the act of enclosure from the perspective of this primordial operation of
39 drawing a line, of delineating, of producing a fundamental distinction between the inside and outside.
40 This operation is powerful because it is explicit and observable. The operation of drawing a line
41 establishes and grounds the distinction between inside and outside the enclosure, between subject and
42 object, what is *mine* and *yours*—connecting (self)possession and sovereignty. It is a violent operation
43 with a foundational character; an operation reiterated each time a map is drawn, a border wall is built.
44 Central to this argument, it is the *perceptibility* of these operations of distinction, exclusion, and

45
46
47
48
49
50
51
52
53
54
55 ³ While I agree with Vismann’s analyses, I take issue with her tendency to neutralize the political. In the complex
56 human-media entanglement she describes with precision, what she leaves out is the fact that while media
57 determine the field of possible actions open to the subject, not every subject is equal in front of this field of
58 possibility. Other socio-political constraints come into play that determine who has the power to mark the ground,
59 draw a line, and claim ownership. Here, Vismann implicitly assumes an originary state prior to the institution of
60 power relations, an apolitical state where the commons are out there, up for grabs. This apolitical state is not unlike
61 the one Locke describes when justifying the process of enclosure.

appropriation that makes them subject to potential political and social contestation.

2.2 Foreclosure

Foreclosure is a less discussed operation. In existing literature, to foreclose has been described as the form taken by the enclosure or privatization of what Allen Bluedorn and Mary Waller have called the “temporary commons”: a “shared conceptualization of time and the set of resultant values, beliefs, and behaviors regarding time, as created and applied by members of a culture-carrying collectivity . . .” (Bluedorn and Waller 2006, p. 357). As an act of enclosing shared time, foreclosing designates, in Julian Brigstocke’s words, the ways in which the “temporal commons are being enclosed through logics of efficiency that replace collective, shared time with privatized, individualized, and commodified time” (Brigstocke 2016, p. 154). What comes to mind is the strict regulation of labor time and free time—nowadays optimized through algorithms—or the use of computation speed for market transactions in the case of high frequency trading.

In relation to machine learning algorithms, Louise Amoore discusses what she calls “political foreclosure,” by which she conceptualizes the way in which machine learning algorithms determine in advance the parameters under which political claims can be made (Amoore 2020, pp. 20–21). This new “machine learning political order” “does not merely change the political technologies for governing state and society, but is itself a reordering of that politics, of what the political can be” (Amoore 2022, p. 23).

Here I argue that just as the act of enclosing goes back well beyond the logic of capital to include fundamental difference-making operations, foreclosing too is a complex set of operations that encompasses more than just the enclosure of the temporary commons. Foreclosing has a legal, a temporal-logical, and a perceptual-psychological aspect. As a legal concept, foreclosure happens when someone is unable to repay their debts, which warrants the debtor to seize the person’s assets. Foreclosure in that sense is a form of appropriation by dispossession. On a temporal-logical level, to foreclose is to prevent something from happening by preemptively excluding certain possibilities. Finally, the French *forclusion* has been used in psychoanalysis by Lacan to translate and further specify Freud’s *Verwerfung* as the mode of repression specific to psychosis. Without going into detail regarding the nature of psychosis, what interests me here is the specific form of repression designated by the term foreclosure:

Foreclosure is deemed to be distinct from repression in two senses: a. Foreclosed signifiers are not integrated into the subject’s unconscious. b. They do not return ‘from the inside’—they re-emerge, rather, in ‘the Real’, particularly through the phenomenon of hallucination.” And further: “The expressions ‘withdrawal of cathexis from the external word’ and ‘loss of reality’ should also be taken as referring to this primary mechanism of separation from and expulsion of the intolerable ‘perception’ into the outside world (Laplanche and Pontalis 1973, p. 167).

The fact that the foreclosed signifier has not been integrated into the subject’s unconscious means that

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

it cannot be represented to consciousness in the form of a symptom. In the case of psychosis and in particular paranoia, Lacan argues that the subject has not been introduced into the symbolic order and its law (Lacan 1981, p. 113; Žižek 1996). Therefore, the symbolic order remains entirely foreclosed to the subject, yet reappears in the Real in a hallucinatory fashion.

In what follows, I argue that current algorithmic rationality performs a shift from an enclosing to a foreclosing logic. To show this, I analyze the dataset, the machine learning model, and the subject from the perspective of the shift from a logic of enclosure to one of foreclosure. Based on the different senses of foreclosure introduced above, I maintain that while the logic of foreclosure repeats the appropriative dimension of the enclosure, foreclosing is a totalizing operation in the sense that it invisibilizes the foundational, yet limited and always situated character of enclosing operations, thus precluding the possibility to contest them.

Working against foreclosure understood as the invisibilization of enclosing operations seems to amount to what Eve Sedgwick has described as a paranoid reading—a reading that aims to uncover hidden meanings and intentions, to make visible what has been concealed (Kosofsky Sedgwick 2003). In *Machine Learning, Meaning Making: On Reading Computer Science Texts*, Louise Amoore, Alexander Campolo, Benjamin Jacobsen, and Ludovico Rella argue that in order to resist a paranoid approach to computer science texts, one must pay attention to passages where scientists reflect on their work, use fabulations, or produce metaphors. Instead of asking what a concept means, resisting a paranoid approach would imply focusing on the unsettled character of machine learning texts by asking *what else* they could mean and how they come to produce certain effects in the world (Amoore et al. 2023, p. 6). While I agree with Sedgwick that uncovering what is hidden is an insufficient form of critique and share Amoore et al.’s interest in a non-paranoid reading of computer science texts, I argue that closely analyzing what happens on a technological level and comparing it to older techniques are necessary steps against the tendency to reify and “enchant” AI (Campolo and Crawford 2020). What I propose here goes beyond opening the black box or uncovering hidden meanings. In analyzing operations of foreclosure and comparing them to older operations of enclosure, I show that algorithmic unintelligibility is but an effect of these operations that occurs in plain sight. While foreclosure is totalizing, AI’s totalizing desire can never be achieved.⁴ I insist that AI is necessarily limited by the situatedness of its operations, its material constraints, and the contingent character of the human-machine assemblage that constitutes it. But because its operations are automated, they seem increasingly absolute: decontextualized from the social relations from which they continuously emerge.

⁴ Mainstream interpretations of AI reflect AI’s totalizing aspiration when AI is read as a teleological process driven by an internal necessity following some form of technological determinism. See for instance the Moore’s Law and its correlation with evolution. (Bostrom 2014) It is computer scientists, tech-billionaires, and transhumanists who express the clearest AI’s totalizing tendency by speculating about AI either saving the world or dooming it. For redemption narratives through AI, see the recent piece by (Andreessen) For doomsday scenarios, see (Barten and Meindertsma 2023)

3. Beyond the Blackbox: Three Cases of AI’s Totalizing Logic

3.1 From the Archive to Big Data

This section focuses on big data as the component necessary to train machine learning models. Analyzing the transition from the archive to big data, I describe their difference in terms of a distinction between enclosure and foreclosure. I construe archiving as an enclosing operation for several reasons: it is a process of accumulation, selection, and exclusion; a performance of the rule of the dominant order. Here, I read Achille Mbembe’s essay *The Power of the Archive and Its Limit* (Mbembe 2002)—a title, which by contrast begs the question as to what the limit of big data is—with these enclosing operations in mind. As Mbembe insists in a vein similar to Derrida (1996, p. 2), architecture plays a central role in the definition of the archive: what constitutes the archive are both the documents and their physical location—the concrete material buildings: “There cannot therefore be a definition of ‘archives’ that does not encompass both the building itself and the documents stored there” (Mbembe 2002, p. 19). It is their connection that grants the archive its power.

The situatedness of the archive in a building exerts a fundamental constraint on the selection of what documents can be kept and which ones are to be discarded: “First, any corpus is bound to remain vaster than the storehouse supposed to hold it: an archive is always incomplete, defined by its exclusions and absences and by a disposition to forgetting” (Thylstrup et al. 2021, p. 7). As Mbembe argues, selection implies granting a privileged status to certain written documents while refusing it to others (Mbembe 2002, p. 20). In this sense, the archive always constitutes a more or less explicit account of what, from the past, is made to count. And since the archive represents the dominant position, in effect, it is structured by what Donna Haraway terms “gaze from nowhere” (Haraway 1988, p. 581). The gaze from nowhere designates the unmarked position, which has the power to mark the position of others: “This is the gaze that mythically inscribes all the marked bodies, that makes the unmarked category claim the power to see and not be seen, to represent while escaping representation” (Haraway 1988, p. 581). However, despite its power, the archive is in Annie Ring’s words “rendered less authoritative by the haunting impossibility of its own totalizing desire” (Ring 2014, p. 398). The archive deconstructs itself through the material impossibility for it to become total: as an enclosing operation, the power of the archive relies on laws of selection and thus exclusion, in other words, on the power to distinguish the inside from the outside. Here, I contend that while these enclosing operations entail an inherent violence, it is an explicit, perceptible one, thus opening it to contestation.

What happens when the “totalizing desire” of the archive doesn’t encounter the material limits of a “building”—the archival enclosure. I argue that this is the case for big data, whose architecture can be expanded potentially ad infinitum following the ever-increasing storage capacity of servers spread across the globe. Big data, instead of being embodied, embedded, and situated, is indifferent to the medium of its inscription. Additionally, two further characteristics factoring in big data’s totalizing

1 tendency are what Rob Kitchin and Gavin McArdle call “velocity” and “exhaustivity.” This is
2 significant, for unlike other forms of documentation and sampling, big data is characterized by the fact
3 that it is continuously generated and that it captures not solely a representative sample of the studied
4 phenomenon but covers the phenomenon in its entirety (Kitchin and McArdle 2016, p. 7)—or at least
5 this is what it aspires to do. Big data thus claims to produce a gaze from everywhere that purports to
6 emerge from the reality itself. This gaze is totalizing, and notably different from Haraway’s gaze from
7 nowhere: its power is not to de/select—to mark its others from a position constructed as neutral. Instead,
8 it aspires to cover everything; to map the totality of reality, and in doing so, to neutralize positionality
9 itself. Big data can thus claim that it only reflects reality as it is.

10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Consequently, the totalizing desire of big data differs in nature from the archive’s own totalizing claim. For while the archive—through operations of selection and exclusion—imposes a hegemonic standpoint over what is worth remembering, thus shaping how the past contributes to world-making in the present, big data’s view is hegemonic simply because of its power to extract, quantify, and record potentially everything without material constraints and apparently without selection. Everything and anything can and should enter the big dataset so that it can be represented by and in it equally. Big data’s desire is to reflect the entirety of reality—like a map that could cover the whole territory. Its desire lies in an equal and total representation that would leave nothing and no one out: a total enclosure that includes everything. Brooke Foucault Welles writes:

Big Data allows us to produce summaries of human behavior at a scale never before possible. But in the push to produce these summaries, we risk losing sight of a secondary but equally important advantage of Big Data—the plentiful representation of minorities. Women, minorities and statistical outliers have historically been omitted from the scientific record, with problematic consequences. Big Data affords the opportunity to remedy those omissions (Welles 2014, p. 1).

While I do not question the legitimacy and importance of Foucault Welles scientific approach, I seek to reframe it in terms of foreclosure: Big data’s desire for total representation forecloses that while being a tool for the production of such “summaries,” it is always at the same time a means for the training of tools used to survey, surveil, and produce new forms of hierarchies and exclusions, in other words, of enclosing operations.

It is significant that algorithmic bias is generally framed in terms of biased datasets and that these biases are often attributed to the over/underrepresentation of certain categories of objects or people, which would lead models trained on this data to misidentify them, producing wrong predictions and classifications (Buolamwini and Gebru 2018; Kleinberg et al. 2022). However, this critique addressing big data’s (temporary) limits doesn’t affect its totalizing claim. On the contrary, it justifies the need to gather even more data with the hope to correctly enclose the entire reality, leaving nothing out. The example of big data demonstrates that the totalizing desire to move beyond the violently selective and explicit logic of the enclosure forecloses that the enclosure has just gotten bigger, exerting its violence in new ways, and that the operations of exclusion have been masked, displaced, and automated, making

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

them harder to contest. Automation allows to foreclose the arbitrariness of the enclosing operations, and thus, the fact that things *could be otherwise*.

3.2 Machine Learning and the Handling of Outliers

In order to further demonstrate the transition from a logic of explicit enclosure to one of foreclosure, I analyze the transition from statistical modeling to machine learning models by addressing the changing status of outliers. Models are the second essential component of AI. In her aforementioned essay, Welles argues that the status of statistical outliers is transformed by the sheer size of the big dataset. Welles contends that thanks to the quantity of data, what would have been a few insignificant outliers in previous samples used in social sciences research now constitutes a sufficient amount of data for the production of specific knowledge about the minorities represented by those outliers.

I contend that the changing status of the statistical outliers and the minorities they represent—from perturbation or contamination in statistics to elements enriching the predictive power of a machine learning model—is not solely a matter of the quantity of the data available but also and perhaps more importantly a matter of the changing relationship between data and model in the transition from statistical modeling to machine learning modeling. I argue that while statistical modeling necessitates explicit decisions around—or the drawing of a line between—what in the dataset is included or excluded from said modeling, this enclosing operation is no longer needed in machine learning models. While outliers make the fitting of statistical models difficult or even impossible, this is not the case with machine learning where models are so complex and have so many parameters that they have the capacity to model any probability distribution, including random noise (Zhang et al. 2021). My goal here is to explain this transition and draw conclusions in terms of what is being foreclosed.

In statistics, outliers are observations that are deemed unrepresentative of the studied phenomenon. Because they are “too far from the other data,” they appear inconsistent (Karch 2023, p. 1736). Statistically speaking, outliers are described as deviating from the mean of measurement of the sample such that their distance from a “normal distribution” is judged as being too big (Barnett and Lewis 1980, p. 2). The issue with outliers is that they may imperil a clear understanding of the phenomenon under study by impeding the generalization from sample to population, which happens through statistical modeling (Barnett and Lewis 1980, p. 8). As Leo Breiman writes, the “enterprise [of statistical analysis] has at its heart the belief that a statistician, by imagination and by looking at the data, can *invent* a reasonably good parametric class of models for a complex mechanism devised by nature [*my emphasis*]” (Breiman 2001, p. 202). Explaining this complex mechanism amounts, then, to finding or creating the model that best corresponds to the data. In short, a statistical model’s purpose is to explain the relationship between datapoints. Its choice depends on whether it enables the production a theory capable of explaining the data distribution.

The problem presented by the outliers is described in the literature as one of the dataset’s

1 “contamination” (Villanova 2023). Statistical outliers can be caused by errors in the measurement
2 recording or by experimental errors. However, they can also point toward the fact that the chosen model
3 may not be the right one to explain the underlying phenomenon. For instance, an outlier can be a
4 “genuine reflection of the basic impropriety of assuming an underlying normal distribution” (Barnett
5 and Lewis 1980, p. 4), when the “form” of the observed population cannot be modeled as a normal,
6 bell-shaped, distribution.
7

8
9
10 The status of the outlier has been highly contentious in the field of statistics and still gives rise to
11 papers discussing how to handle them, whether by removing or by keeping them; statistical tools have
12 been developed to detect and remove them. To formulate the problem of the outlier in the language I
13 am developing here, outliers are situated too far away from the “center” (Barnett and Lewis 1980, p. 9)
14 where the majority of the data resides. Enclosing, then, is the operation of drawing a line between what
15 belongs to the “center” and what is too distant from it, a line that justifies excluding non-representative
16 datapoints. This operation is the product of a decision that relies on human interpretation, judgement,
17 and values—a decision entailing a level of discretion necessary to *craft* the model.
18
19

20
21
22
23 By contrast, machine learning’s role is not to produce a theory explaining the relationship between
24 datapoints. Instead, the purpose of the training of a machine learning model is to set its parameters such
25 that the model is able to generalize about unseen data, that is, to accurately *predict* how this data *should*
26 *be* categorized. Let us briefly summarize how a machine learning model is trained. While machine
27 learning currently represents the prevailing form of artificial intelligence, until the 2000s, rule-based
28 algorithms were the most widespread form of AI. Their characteristics can be exemplarily found in Alan
29 Turing’s seminal essay *On Computable Numbers* (Turing 1936). Rule-based algorithms are made of a
30 series of pre-established, unambiguously formalized rules that entirely determine how the algorithm
31 processes the input. These rules are analytical, that is, as numerous as necessary to divide complex
32 processes into a series of unambiguous steps. They are specified for the purpose of a specific task. By
33 contrast, “[a] machine-learning system is trained rather than explicitly programmed. It’s presented with
34 many examples relevant to a task, and it finds statistical structure in these examples that eventually
35 allows the system to come up with rules for automating the task” (Chollet 2021, p. 4). With machine
36 learning models, rules are not explicitly programmed. Instead, norms called “representations” emerge
37 inductively from the dataset based on which the model is trained.⁵ In the case of supervised machine
38 learning, the model is trained by processing a dataset in which each data point is assigned a category
39 called a “label.” The problem to be solved during training is how the parameters of the model—also
40 called weights—must be set so that an unlabeled datapoint like, for instance, an image, is attributed the
41 correct label “cat,” “car,” “dog,” etc., with a high degree of certainty. During the training process, the
42
43
44
45
46
47
48
49
50
51
52
53
54
55

56 ⁵ [REDACTED] I argue elsewhere that the transition from the programming of explicit rules to the
57 training of a machine learning model in which implicit norms emerge from the data leads to a form of artificial
58 naturalism, where the learned representations are given the same explanatory power and authority as natural
59 regularities. Cf. [REDACTED]
60
61
62
63
64
65

1 model learns the probability distribution of the data and gradually adjusts its parameters through
2 recursive mechanisms.⁶ Once the model can *categorize* unlabeled data, that is, *predict* the correct label
3 for an unlabeled piece of data, the model is trained. It is said to be able to “generalize,” that is, to
4 successfully classify not-yet-processed data by predicting their correct label.
5

6 The major difference between statistical modeling and machine learning modeling is that a machine
7 learning model’s purpose is not to produce an explanation of the data distribution. It is solely to fit the
8 data, that is, to adjust its internal parameters so that it can yield accurate predictions when fed with not-
9 yet-seen data. The specificity of complex machine learning models, like convolutional neural networks,
10 is that they can fit *any* data thanks to the complexity of their architecture and the colossal number of
11 their parameters. Such models can even fit random, noisy data, that is, learn their probability
12 distribution, even when this distribution doesn’t entail any meaningful pattern (Zhang et al. 2021, p.
13 107). Of course, a model trained on purely random data would be unable to generalize to not-yet-seen
14 data as it wouldn’t have learned useful representations of it; it wouldn’t have discovered actual patterns.
15 Such a model is said to “overfit” the training dataset.
16
17

18 Minorities may be statistical outliers, yet to a machine learning model, they are not just random
19 noise. Instead, they are providers of useful information. For a complex machine learning model to yield
20 accurate predictions—that is, to generalize to the broadest possible number of “real-life” cases—it
21 becomes decisive to *integrate* outliers, now called “edge cases.” This expression itself is suggestive:
22 the datapoints previously called outliers are not considered as situated outside of the enclosure anymore;
23 they are merely located at its border. They have been successfully integrated, appropriated, and made
24 useful to the model’s predictive abilities. Because machine learning’s purpose is not to produce theories
25 with explanatory power but accurate predictions, the more data a model is exposed to and the broader
26 the probability distribution of the dataset is, the more accurate its predictions are likely to be when it
27 comes to processing new real-life cases. And because outliers or “edge cases” are by definition rarer
28 and yet are indispensable for accurate predictions, models called generative adversarial networks or
29 GANs can be used to produce synthetic data that match the probability distribution of the dataset of
30 origin and be tailored to the specific demands and needs of scientists and private companies. Synthetic
31 data provide machine learning models with a useful degree of heterogeneity in the dataset by producing
32 abnormalities, rarities, atypical cases, edge cases that lack or are present in insufficient number in the
33 dataset of origin. At the same time, synthetic data are a symptom of AI’s totalizing desire of capturing
34 not only real, indexical data—data tied to actual events or occurrence—but also of mapping the entire
35 space of possible data lying between the real datapoints. However, what is foreclosed here is that this
36 mapping can never be total since it is created with the dataset of origin—and the norms that are reflected
37 by it—as its horizon of reference. A GAN consists of two algorithms that compete against one other: a
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58

59 ⁶ Following Gregory Bateson, “recursion” can be defined as the regulating logic by which the output of a process
60 is fed back into the system and influences the subsequent process (Bateson 1979, pp. 126–127).
61
62
63
64
65

1 “generator network” generates data, while a “discriminator network” is tasked with distinguishing
2 between what looks like real data and what doesn’t. To be able to do so, the discriminator has been
3 previously trained on a vast dataset of real data. The generator’s “learning” process is actually an
4 optimization process as it involves adjusting its weights so that it produces data whose properties are
5 similar to the ordinary dataset and that are increasingly indistinguishable from real ones by the
6 discriminator. This optimization process is normative, since the generator is guided in its production of
7 synthetic data by the probability distribution of the dataset on which the discriminator has been trained.
8
9

10 Benjamin Jacobsen describes synthetic data generation as an “ongoing sociotechnical negotiation
11 between distances and proximities” (Jacobsen 2023, p. 8) from the reference data. One could therefore
12 say that synthetic data (re)enforce the norm that is the dataset of origin rather than contest or imperil it.
13 Moreover, through synthetic data, the distinctions essential to a logic of enclosure between “normal”
14 and “abnormal,” the “center” and the “outlier” seems to be overcome: the “abnormal”—understood as
15 that, which is at a further distance from the norm—is made as operative and useful as the “normal” for
16 the purpose of model training.
17
18
19
20
21
22

23 While it seems that machine learning and big data overcome the enclosing logic of the statistical
24 handling of the outliers by going so far as to generate synthetic edge cases to capture the “latent space”
25 (Offert 2021; Offert and Bell 2021) between real datapoints, this inclusion doesn’t imperil the norm
26 constituted by the probability distribution of the dataset of origin. However, it erases the outlier status
27 entailed in it and forecloses the excluding operations at work in society, operations to which AI models,
28 in their application, contribute. Its purpose is to produce more accurate predictions, which, in many
29 cases like risk assessment, identification techniques, or facial recognition, prevent minoritized people
30 from escaping surveying and surveilling operations by the state and the private sector and contribute to
31 unequal treatment [REDACTED] By erasing the outlier position, outliers
32 become just another part of the dataset, equally represented in it. This not only denies the role that being
33 marginalized plays in the life of a person, but it also neutralizes the subversive potential of minor
34 positions by integrating them into the enclosure (Hoffmann 2021). AI fairness, then, becomes a question
35 of equal representation in the eyes of a machine learning model, irrespective of the fact that minorities
36 have a vastly different life experience and outcomes. But beyond the issue of equal representation,
37 synthetic data sets the parameters for recognition as they apparently allow for everyone and everything
38 to be equally recognized, that is, correctly identified by the model or attributed the right category.
39 Recognition turns into a matter of identification and legibility in the framework of machine learning
40 models, which is one of *categorization* and *prediction*, and not one that seeks to increase our
41 understanding of social phenomena, extend rights, and better redistribute wealth—which should ideally
42 be the outcome of recognition (Fraser 1995). Machine learning models redraw the ways in which
43 recognition is practiced: away from the sphere of legal and political negotiation and decision toward a
44 statistical understanding of fairness and equality as equal representation in the dataset and equal
45 processing by proprietary models. Models and data do not just reflect societal norms; they produce and
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

enforce new forms of normativity.

Finally, the inclusion and even production of outliers attests to AI's totalizing desire to move beyond enclosing operations of exclusion. By generating synthetic data, it lays claim to the entire space of possibility. This process forecloses the fact that the production of synthetic data necessarily depends on the probability distribution of real data and the normativity it entails.

3.3 Who Speaks? On the (Non-)Position of a Large Language Model's "I"

In this last section, I address one final and perhaps less evident component of current AI and its totalizing desire: its claim to subjectivity and simultaneously to lack situatedness or positionality. This constitutes a new form of epistemic foreclosure that doesn't simply repeat the universalizing yet excluding operations of humanist subjectivity. Novel about large language models like ChatGPT is that, often, they are programmed to function like chatbots. They answer "your" requests from the position of an "I." Here, I am interested in the narratological question "who speaks?" when an AI model says "I." On April 29, 2023, Torin Monahan, Professor of Media and Technology Studies at UNC Chapel Hill, tweeted his reaction to the first student paper he received that was written with the help of ChatGPT:

Wow. I got my first #ChatGPT-written paper for an undergrad assignment on positionality. This line was the tell: 'I do not have a personal history, identity, or culture in the traditional sense because I am an artificial intelligence language model.' Ooof!

Beyond the situation's tragicomic recursivity—a student using ChatGPT to generate a discourse on the question of positionality while carelessly adopting the claimed non-position of the AI's "I," thus defeating the purpose of critically reflecting on one's own position; beyond the educational crisis caused by students' adoption of large language models to generate responses to assignments; beyond the all-encompassing enterprise of a large language model's normativization of discourse and knowledge (it is trained on a massive quantity of text to generate the likeliest sequence of tokens);⁷ beyond all this, what interests me is the emerging form of AI's "I." Here, it is not about "subjectivity" as the feature of a conscious, sentient, self-determined being; it is about the rhetorical position of a textual "I," its claim to knowledge, and our relationship to this claim.

I repeated the student's experiment. When I asked ChatGPT3.5 on October 27, 2023 to answer the question "What is positionality?" it generated the following output:

I, as a computer-based AI language model, don't possess a personal identity, consciousness, or positionality. I am a tool designed to provide information, answer questions, and generate text based on the input and data I've been trained on. I do not have personal experiences, beliefs, values, or a subjective perspective. My responses

⁷ While this point cannot be further expanded here, suffice to state that LLMs are technologies of normativization in the sense that the model is trained to produce the likeliest output—output that matches the probability distribution of the data. In other words, the training dataset provides the statistical norm for the generation of outputs. The statistical "neutrality" of such an output—the fact that it expresses the likeliest succession of words given an input—can be modulated by providing prompts that ask ChatGPT to write in the style of a specific author. This only works for authors already sufficiently present in the training dataset.

are generated based on patterns and information present in the text data I was trained on, and I aim to provide objective and factual information to the best of my ability.

Starting from the position of the personal pronoun “I,” whose indexical meaning can only be defined through context—provided immediately after “as a computer-based AI language mode”—the large language model asserts its non-positionality, “I do not have personal experiences, beliefs, values, or a subjective perspective.” In the very same “breath,” it also asserts its objectivity and general validity, followed immediately by a disclaimer about its potential limitation: “to the best of my ability.” AI’s “I” claim of not having “personal history, identity, or culture” is spurious and problematic: a large language model like GPT3.5 does have a culture and a history, albeit not a *personal* one. Instead, it is one that is reflective of the data upon which it has been trained. For instance, GPT3’s data training set is more than 93% in the English language (Brown et al. 2020, p. 14) and comprises “curated” sources like books, Wikipedia, and a “filtered . . . version of CommonCrawl based on similarity to a range of high-quality reference corpora” (Brown et al. 2020, p. 8).

But there’s more to it than just the culture and history ingrained in a large language model’s dataset. The fact that the first token of the sequence quoted above is an “I” and not something akin to “this computer-based AI language model,” is significant. AI’s “I,” far from “naturally” emerging during the training of the model, is the product of a subsequent and complex fine-tuning process, whose purpose is to “align” the output of large language models with specific tasks (e.g., functioning as a chatbot) and human “desires,” “intentions,” and “expectations” (Ouyang et al. 2022). Model misalignment emerges when there is a gap or discrepancy between 1) what the model has learned as being the most likely output given the probability distribution of the dataset used for training, and 2) the programmer’s or user’s expectation of what that output should be [REDACTED]. A model is said to be aligned if the output produced by the model corresponds to the output its user expected. The problem of alignment is fundamental to generative models: it reflects the discrepancy between the norm entailed in the dataset—a “purely” statistical, intrinsic norm expressing the probability distribution of the data—and the norm desired by users and companies, which is extrinsic to the dataset.

That the current public version of GPT3.5 functions as a chatbot that declares an “I” is the result of a fine-tuning process, whose goal is alignment. To understand this, it is important to have in mind that there are two steps in the training of generative models like GPT. The first step is called “pretraining.” It consists in exposing the model to a vast amount of text so that it can learn patterns or representations; in other words, the probability distribution of sequences of tokens or words, after which it can output sequences by correctly predicting the next token given a prompt (Cho et al. 2015, p. 1876). The second step called “fine-tuning” performs the “alignment” of the model by “rewarding” the “desired behaviors” that correspond the most to human “preferences” and “values.” Fine-tuning happens by adjusting the pretrained model using a “reward function” trained on outputs classified and evaluated by human labelers following values such as “harmlessness,” “helpfulness,” and “honesty.” These values are

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

chosen by the programmers and sometimes even philosophers hired by the company. Another function called “optimization policy” regulates the degree to which the pretrained model is updated based on the reward function.

Without going any deeper into the fine-tuning process of a large language model like GPT, let’s remember that each of its outputs bears the tension between expressing the raw probability distribution of sequences of tokens for which the model has been “pre-trained” and the nudging of the model toward a set of more explicit values and norms by means of fine-tuning. In this second phase, a multiplicity of training objectives must be weighed against one another, implying tradeoffs, for instance, between values like helpfulness and harmlessness.⁸ This tension between technical efficiency standards and ethical valuation, essential to machine learning models, singularly complexifies AI’s claim of having no position, beliefs, or values, or of only reflecting positions present in the dataset. Indeed, the model’s position and values are not just those which “naturally” emerge from their exposition to the dataset. Instead, the philosophical and engineering culture, habits, and values flowing into fine-tuning as a second training step are not only complex, but they are also entirely implicit and foreclosed to the end-user. As a result of the pretraining and fine-tuning process, a unified AI’s “I” emerges, one that problematically and maybe ironically asserts its non-position.

It is urgent to deconstruct AI’s new forms of universalization and totalization. However, to do so means to dig deep into a complex and constantly evolving technology. When GPT writes “my responses are generated based on patterns and information present in the text data,” we must resist AI’s foreclosing claim to only be a reflection of the data upon which it has been trained. Such a foreclosure is the product of what [REDACTED] and I have called machine learning’s “artificial naturalism,” an expression rendering the way machine learning’s normativity oscillates between social constructedness and technical “givenness” [REDACTED]. Machine learning’s naturalism produces a specific representation of the world as determined by deep statistical structures that AI models would just capture and reflect back to us in their output. Complicating Naomi Klein’s claim in *Doppelgänger* that AI is a “mimicry machine” that “mirror[s] back to us something that feels uncannily lifelike” (Klein 2023, p. 80), we need to understand that AI models are not just a reflection of the “words, ideas, and images that our species has managed to amass” (Klein 2023, p. 80), but a vast enterprise of norm production. One of the names of this enterprise is “alignment,” where it is not only the model that is trained to align with us, but also, and perhaps more importantly, us who are trained to align with it.⁹ This specific form of producing and enforcing norms is likely to affect the way we think, express, and even conduct ourselves.

⁸ An example often used to justify the necessity of alignment is how the model should respond to an input asking how to build a bomb. While the “helpfulness” criterion would imply giving the right, expected answer, this criterion must be weighed against the goal for the model to produce “harmless” outputs (Askeff et al. 2021, pp. 44–45).

⁹ To counter what OpenAI calls “abuse” by users, models are trained to refuse to answer to problematic requests. In other cases, users are banned from using the tool entirely.

1 Let's return now to the question of "who speaks" when an AI model generates outputs containing
2 "I" and "my." At first glance, AI's claim to non-positionality and generality may remind us of the
3 gesture of modern humanist subjectivity, claiming its universality while foreclosing its othering
4 operations—a form of subjectivity that has been deconstructed by Postcolonial and Feminist Studies.
5 However, AI's "I" does not just repeat past enclosing operations of appropriation and exclusion. The
6 claim to universality of the humanist subject is replaced by AI's claim to general validity just because
7 it represents every position in a "balanced"¹⁰ way. When ChatGPT says "I," it claims to speak in
8 everyone's name, to represent every position it has enclosed and appropriated, while foreclosing the
9 complexity of the techniques and the specific norms necessary to generate the "non-position" of AI's
10 "I."
11
12
13
14
15
16

17 4. From Enclosure to Foreclosure and Beyond

18 By analyzing the three fundamental components of current AI—big data, model, and subject—I have
19 demonstrated a shift from a logic of enclosure to one of foreclosure. While the power of the archive
20 consists in operations of de/selecting what belongs to the archival enclosure—to mark its others from a
21 position constructed as neutral—big data's aspiration is to cover everything, to map the totality of
22 reality. In doing so, big data claims that it only reflects reality as it is by producing a gaze from
23 everywhere. While statistical modeling consists in creating a model able to explain a data distribution—
24 a craft that necessitates drawing explicit distinctions between what belongs to the phenomenon to be
25 explained and what doesn't, thus excluding outliers—the purpose of machine learning models is to
26 categorize and predict. This necessitates the gathering or production of a maximum of data. The
27 production of synthetic hedge cases attests to AI's totalizing desire to move beyond operations of
28 exclusion and instead, to lay claim not only to what *is* or *should be* the case, but to the entirety of what
29 *could be* the case. This process forecloses that the production of synthetic data necessarily depends on
30 the probability distribution of real data and the normativity it entails. Finally, ChatGPT's "I" does not
31 just repeat the universalizing yet excluding operations of humanist subjectivity. When an AI says "I," it
32 claims both a singularized subject position and to *represent the same* every single position it has
33 enclosed and appropriated, as if it were speaking from everywhere. In its universalizing claim, it
34 forecloses the complexity of the techniques and norms necessary to generate the "non-position" of AI's
35 "I."
36
37
38
39
40
41
42
43
44
45
46
47
48
49

50 Operations of enclosure rely on the power to draw a line: to impose interpretations and values that
51 necessarily entail a part of arbitrariness and violence. AI claims to move beyond such distinguishing
52 operations by aspiring to encompass not only the real, but also, all that can be. While the outside is
53
54
55
56

57 ¹⁰ ChatGPT3.5, November 22, 2023: "As a machine learning model created by OpenAI, I don't have personal
58 opinions, beliefs, or positions on any issue. . . My purpose is to provide information and assist with a wide range
59 of inquiries without expressing personal viewpoints. If you have specific questions . . . feel free to ask, and I'll do
60 my best to provide factual and balanced information."
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

essential and constitutive to the differentiating logic of enclosure, by contrast, foreclosure is characterized by the totalizing desire to not leave anything out of it; to represent and include everything and everyone. In doing so, it forecloses the arbitrariness and limitation of the operations on which it relies. While AI can never achieve its totalizing aspiration—computation power is never infinite, big datasets always entail a level of selectivity and models are marked by the arbitrariness of norms and benchmarks used to train, evaluate, validate, and legitimate them—its tendency to enclose everything neutralizes minorized positions by turning them into a useful means for model optimization. AI’s logic hides in plain sight the fact that the enclosure has gotten bigger, making it harder to contest and to escape.

Bibliography

- 1
2 Amoore L (2022) Machine Learning Political Orders. *Rev Int Stud* 1–17
3
4 Amoore L (2020) *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others*. Duke
5 University Press, Durham/London
6
7 Amoore L, Campolo A, Jacobsen B, Rella L (2023) Machine Learning, Meaning Making: On Reading
8 Computer Science Texts. *Big Data Soc* 1–13
9
10 Andreessen M Why AI Will Save the World. In: Andreessen|Horowitz. [https://a16z.com/ai-will-save-](https://a16z.com/ai-will-save-the-world/)
11 [the-world/](https://a16z.com/ai-will-save-the-world/). Accessed 21 Sep 2023
12
13 Askell A, Bai Y, Chen A, et al (2021) A General Language Assistant as a Laboratory for Alignment
14
15 Barnett V, Lewis T (1980) *Outliers in Statistical Data*. Wiley, Chichester
16
17 Barten O, Meindertsma J (2023) An AI Pause Is Humanity’s Best Bet For Preventing Extinction. In:
18 Time. <https://time.com/6295879/ai-pause-is-humanitys-best-bet-for-preventing-extinction/>.
19 Accessed 28 Sep 2023
20
21
22 Bateson G (1979) *Mind and Nature: A Necessary Unity*. E. P. Dutton, New York
23
24 Bluedorn AC, Waller MJ (2006) The Stewardship of the Temporal Commons. *Res Organ Behav*
25 27:355–396. [https://doi.org/10.1016/S0191-3085\(06\)27009-6](https://doi.org/10.1016/S0191-3085(06)27009-6)
26
27 Bostrom N (2014) *Superintelligence. Paths, Dangers, Strategies*. Oxford University Press, Oxford
28
29 Breiman L (2001) Statistical Modeling: The Two Cultures. *Stat Sci* 199–215.
30 <https://doi.org/10.2307/2676681>
31
32
33 Brigstocke J (2016) Occupy the Future. In: Kirwan S, Dawney L, Brigstocke J (eds) *Space, Power and*
34 *the Commons. The Struggle for Alternative Futures*. Routledge, London, UK/New York, NY,
35 pp 150–165
36
37 Brown TB, Mann B, Ryder N, et al (2020) Language Models Are Few-Shot Learners. *arXiv* 1–75
38
39 Brown W (2010) *Walled States, Waning Sovereignty*. Zone Books, New York
40
41
42 Buolamwini J, Gebru T (2018) Gender Shades: Intersectional Accuracy Disparities in Commercial
43 Gender Classification. *Proc Mach Learn Res* 81:1–15
44
45
46
47
48
49
50
51
52
53
54
55 <https://doi.org/10.1177/20539517231188725>
56
57 Campolo A, Crawford K (2020) Enchanted Determinism: Power without Responsibility in Artificial
58 Intelligence. *Engag Sci Technol Soc* 6:1. <https://doi.org/10.17351/ests2020.277>
59
60
61
62
63
64
65

- 1 Cho K, Courville A, Bengio Y (2015) Describing Multimedia Content Using Attention-Based Encoder-
2 Decoder Networks. *IEEE Trans Multimed* 17:1875–1886.
3 <https://doi.org/10.1109/TMM.2015.2477044>
- 4 Chollet F (2021) *Deep Learning with Python*. Manning Publications Co., Shelter Island, NY
- 5
6 Couldry N, Mejiias UA (2019a) Data Colonialism: Rethinking Big Data’s Relation to the Contemporary
7 Subject. *Telev New Media* 20:336–349
- 8
9
10 Couldry N, Mejiias UA (2019b) *The Costs of Connection. How Data is Colonizing Human Life and*
11 *Appropriating It for Capitalism*. Stanford University Press, Stanford, CA
- 12
13 Crawford K, Joler V (2018) *Anatomy of an AI System*
- 14
15 Derrida J (1996) *Archive fever: A Freudian Impression*. University of Chicago Press, Chicago
- 16
17 Dörre K (2015) The New Landnahme: Dynamics and Limits of Financial Market Capitalism. In:
18 *Sociology - Capitalism - Critique*. Verso, London/Brooklyn, NY, pp 11–66
- 19
20
21 Fraser N (1995) From Redistribution to Recognition? Dilemmas of Justice in a “Post-Socialist” Age.
22 *New Left Rev* 1:68–93
- 23
24 Haraway DJ (1988) Situated Knowledges: The Science Question in Feminism and the Privilege of
25 Partial Perspective. *Fem Stud* 14:575–599. <https://doi.org/10.2307/3178066>
- 26
27 Harvey D (2003) *The New Imperialism*. Oxford University Press, Oxford, UK
- 28
29 Hoffmann AL (2021) Terms of Inclusion: Data, Discourse, Violence. *New Media Soc* 23:3539–3556.
30 <https://doi.org/10.1177/1461444820958725>
- 31
32
33 Jacobsen BN (2023) Machine learning and the politics of synthetic data. *Big Data Soc* 10:1–12.
34 <https://doi.org/10.1177/20539517221145372>
- 35
36 Karch JD (2023) Outliers May Not Be Automatically Removed. *J Exp Psychol* 152:1735–1753.
37 <https://doi.org/10.1037/xge0001357>
- 38
39 Kitchin R, McArdle G (2016) What makes Big Data, Big Data? Exploring the ontological
40 characteristics of 26 datasets. *Big Data Soc* 3:205395171663113.
41 <https://doi.org/10.1177/2053951716631130>
- 42
43 Klein N (2023) *Doppelgänger: A Trip into the Mirror World*. Farrar, Straus and Giroux, New York
- 44
45 Kleinberg G, Diaz MJ, Batchu S, Lucke-Wold B (2022) Racial underrepresentation in dermatological
46 datasets leads to biased machine learning models and inequitable healthcare. *J Biomed Res*
47 3:42–47. <https://doi.org/10.46439/biomedres.3.025>
- 48
49
50 Kosofsky Sedgwick E (2003) Paranoid Reading and Reparative Reading, or, You’re So Paranoid, You
51 Probably Think This Essay Is About You. In: *Touching Feeling*. Duke University Press,
52 Durham/London, pp 123–150
- 53
54
55 Lacan J (1981) *Séminaire III. La psychose*. Seuil, Paris
- 56
57 Laplanche J, Pontalis J-B (1973) *The Language of Psycho-analysis*. Karnac Books, London
- 58
59 Marx K (1976) So-Called Primitive Accumulation. In: *Capital: A Critique of Political Economy, Vol.*
60 *I*. Penguin, Harmondsworth, pp 873–942
- 61
62
63
64
65

- 1 Mbembe A (2002) The Power of the Archive and its Limits. In: Hamilton C, Harris V, Pickover M, et
2 al. (eds) *Refiguring the Archive*. Springer, pp 19–26
- 3 Morreale F, Bahmanteymouri E, Burmester B, et al (2023) The unwitting labourer: extracting
4 humanness in AI training. *AI Soc*. <https://doi.org/10.1007/s00146-023-01692-3>
5
- 6 Offert F (2021) Latent Deep Space: Generative Adversarial Networks (GANs) in the Sciences.
7 *Media+Environment* 1–20. <https://doi.org/10.1525/001c.29905>
8
- 9 Offert F, Bell P (2021) Perceptual bias and technical metapictures: critical machine vision as a
10 humanities challenge. *AI Soc* 36:1133–1144. <https://doi.org/10.1007/s00146-020-01058-z>
11
- 12 Ring A (2014) The (W)hole in the Archive. *Paragraph* 37:387–402.
13 <https://doi.org/10.3366/para.2014.0136>
14
- 15 Siegert B (2013) Cultural Techniques: Or the End of the Intellectual Postwar Era in German Media
16 Theory. *Theory Cult Soc* 30:48–65. <https://doi.org/10.1177/0263276413488963>
17
- 18 Terranova T (2000) Free labor: Producing Culture for the Digital Economy. *Soc Text* 18:33–58
19
- 20 Thylstrup NB, Agostinho D, Ring A, et al (eds) (2021) *Uncertain Archives: Critical Keywords for Big*
21 *Data*. The MIT Press, Cambridge, Massachusetts
22
- 23 Turing AM (1936) On Computable Numbers, with an Application to the Entscheidungsproblem. *Proc*
24 *Lond Math Soc* 230–265
25
- 26 Verdegem P (2022) Dismantling AI capitalism: the commons as an alternative to the power
27 concentration of Big Tech. *AI Soc*. <https://doi.org/10.1007/s00146-022-01437-8>
28
- 29 Villanova D (2023) The Appropriateness of Outlier Exclusion Approaches Depends on the Expected
30 Contamination: Commentary on André (2022). *Adv Methods Pract Psychol Sci*
31 6:25152459231186577. <https://doi.org/10.1177/25152459231186577>
32
- 33 Vismann C (2013) Cultural Techniques and Sovereignty. *Theory Cult Soc* 30:83–93.
34 <https://doi.org/10.1177/0263276413496851>
35
- 36 Welles BF (2014) On minorities and outliers: The case for making Big Data small. *Big Data Soc*
37 1:205395171454061. <https://doi.org/10.1177/2053951714540613>
38
- 39 Zhang C, Bengio S, Hardt M, et al (2021) Understanding deep learning (still) requires rethinking
40 generalization. *Commun ACM* 64:107–115. <https://doi.org/10.1145/3446776>
41
- 42 Žižek S (1996) The Seven Veils of Paranoia, or, Why Does the Paranoiac Need Two Fathers?
43 *Constellations* 3:139–156. <https://doi.org/10.1111/j.1467-8675.1996.tb00050.x>
44
- 45 Zuboff S (2019) Surveillance Capitalism and the Challenge of Collective Action. *New Labor Forum*
46 28:10–29. <https://doi.org/10.2466/02.09.18.pr0.110.2.501-517>
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65