PLAY IT AGAIN

A Genealogy of Machine Learning

Tyler Reigeluth tyler.reigeluth@univ-catholille.fr
Université Catholique de Lille, ETHICS, 0000-0002-9908-1052

Abstract

This article proposes a genealogy of machine learning by contrasting
two normative regimes: play and automation. Each one expresses
differing, albeit not incompatible values, expectations and objectives
when evaluating human-machine behaviors. Whereas automation
pushes us to consider machines as ideally *working by themselves* and
thus replacing humans, play involves a social and affective
engagement whose outcome is always partially unpredictable. More
specifically, I suggest that play provides a sweet spot for thinking
about performances such as those we are increasingly seeing in
machine learning systems, that combine both rule-following behaviors
and forms of improvisation upon those rules. The genealogical
approach serves to foreground certain aspects of machine learning
which have been foreclosed, the point being not to unveil its hidden
or deeper truth, but to make us rethink what we claim to know of our
relationships with machines and how we might engage with them
differently.

1 INTRODUCTION

Contemporary discourses accompanying the deployment of machine learning tend to fit within
a meta-narrative of automation. Whether heralded as inevitable or criticized as reductive, this
process of automation tends to be considered as concurrent to technology's extension in
general. By drawing on the social and technical history of machine learning, I will suggest that
the concept of play offers at least one other way of framing machine learning's development,
one that could help foster other expectations and evaluations of machine learning
performances. My main claim is that play can account for the relative open-endedness of
machine learning behaviors as well as the forms of interaction they foster with their users
within social practices. While I will draw on some of the material and conclusions developed
elsewhere, the focus of this paper is not to discuss what a theory of learning could, or should,
be for machine learning (Reigeluth & Castelle, 2021), nor how these sociotechnical assemblages
should be qualified ontologically or epistemologically (Grosman & Reigeluth, 2019; Mackenzie,
2017). Rather, it is to explore the normative discourse on the relationships between humans
and machines, on what is expected of machines when they "learn", and what humans are
supposed to do when machines begin to learn. The *genealogy* of machine learning I will provide
problematizes contemporary debates by drawing on certain recurring tropes. I will distinguish
two normative paradigms: automation and play. Each one expresses differing, albeit not
incompatible values, expectations and objectives when evaluating human-machine behaviors.
Whereas automation pushes us to consider machines as ideally *working by themselves* and thus
replacing humans, play involves a social and affective engagement whose outcome is always
partially unpredictable. Technological performances are never just a matter of increasing
efficiency, but also about experimenting with multiple outcomes regardless of whether we
know exactly how they were produced (Dumouchel & Damiano, 2016; Ash, 2015; Pickering,
1995; Flusser, 2004). As the historian of technology, Jonathan Sawday (2007, p. 111) underlines
in his book on the rise of machine culture in the Renaissance, "play and fantasy [...] have
perhaps been far greater elements in the evolution of different forms of technology than is
suggested by the that popular (but wrong-headed) belief that 'necessity is the mother of
invention'." As we shall see, play is both a specific form of "experience" (Caillois, 1967) or
"attitude" (Henriot, 1969) in which we are free to engage and a "social function" (Huizinga,
1951) which primes and subtends everyday practices. More specifically, I underline why play

provides a sweet spot for thinking about performances such as those we are increasingly seeing in machine learning systems, that combine both rule-following behaviors and forms of improvisation upon those rules.

I will first look at how machine learning came to be framed as the ultimate step in the process of automation or how industrial labor extended to learning itself. I will then show how this framing serves as a normative ideal rather than a concrete understanding of how machine behaviors are embedded within social practice. As an exemplification of this embeddedness, I will consider how debates around machine learning biases might be better understood in terms of socially instituted education. This will allow me to turn my attention to the (often neglected) pervasiveness of game-playing throughout the history of machine learning in the very material and situated strands of its development, as well as in the more general epistemological context of game theoretical debates and in the political stakes of war games. These elements, moments and distinctions will provide me with a basis to propose a final juxtaposition between two forms of play – divergent and convergent – which can help determine whether particular machine learning systems close or open unforeseen spaces of interactions. In keeping with my genealogical approach, each argument unfolded in this paper serves to foreground certain aspects of machine learning which have been foreclosed, the point being not to unveil its hidden or deeper truth, but to make us rethink what we claim to know of our relationships with machines and how we might engage with them differently.

2 LABOR AND (MACHINE) LEARNING

In his 1964 book, *Understanding Media,* Marshall McLuhan predicted a world in which learning would become the driving economic activity at the expense of fragmented and mind-numbing labor. Paid learning is already becoming both the dominant employment and the source of new wealth in our society. This is the new role for men in society, whereas the old mechanistic idea of 'jobs,' or fragmented tasks and specialist slots for 'workers', becomes meaningless under automation." (McLuhan, 1964, p. 304) Almost sixty years later this extension of learning and creativity promised by automation seems to have acquired quite a different form than the one envisioned by McLuhan. While actors such as Cisco publish white papers claiming that the "the emergence of new technologies, designed around the principles of creativity and collaboration, are likely to prove central to both 21st century pedagogies and 21st century skills", Amazon Turkers are paid to perform micro "Human Intelligence Tasks", which generally involve mechanically clicking on images to help "intelligent" algorithms along in their learning process. A strange new "geography of intelligence" (Schaffer, 1994) seems to be emerging in which machines are supposed to be liberating human creativity and learning at an unprecedented scale all the while rivaling humans in these very faculties. We are increasingly faced with the prospect that machines are not only good at automating menial and repetitive tasks but may excel at socially and cognitively complex activities such as learning and creative problem-solving. Contemporary proponents of *machine learning* even go so far as contending that we are experiencing a third historical revolution: "The Industrial Revolution automated manual work and the Information Revolution did the same for mental work, but machine learning automates automation" (Domingos, 2015, p. 10). In popular as well as academic discourse, the "automation of learning" is presented both as the next if not final stage of mechanization and industrialization, and as the expression of the "postindustrial" technologies that have come to epitomize a new phase of capitalism. Implicit in this way of framing machine learning is a tension between, on the one hand, the idea that learning itself can be understood as a form of labor or productive activity and, on the other hand, the idea that there are machines capable of heretofore exclusively organic abilities such as learning.

Without diving into classical economic theories, the relationship between labor and automation is deeply rooted in the transformations brought about by the Industrial Revolution (Friedmann,

1946; Zuboff, 1988; Noble, 1984). More recently, machine learning's relationship to labor has gained traction, namely within the "digital labor" literature (Cardon & Casilli, 2015; Crary, 2013; Steiner, 2012; Stiegler, 2015). What McLuhan's passage stages even more accurately and presciently than the great emancipation of labor is the way in which learning and leisure have themselves been put to work through a logic of automation. In its endless extension, the "learning society" blended into "cognitive capitalism" (Moulier-Boutang, 2012) and the "attention economy" (Citton, 2017) whereby user experience, skill, and attention, namely on digital platforms, are harvested and monetized through a general process of "gamification" and "seamless" design. From this perspective, even when a technology's "user" is not a worker, it is still unwittingly working. To help circumscribe this uneasy relationship between labor, automation and, learning we need to briefly retrace how classical "intellectual" activities such as mathematic operations came to be automated, that is turned into mechanical operations analogous to manual labor. The social history of the computer provides a perfect case in point of this mechanization of intellectual labor afforded by a preliminary division of labor and intelligence.

Indeed, it is commonplace to reiterate that the first computers were human agents (Grier, 2007; Mindell, 2002) – often women (Light, 1999; Wise, 2007) – assigned to specific mathematical functions or operations, like cogs within a larger computational machine. A deskilling of computation occurred throughout the late 18th century as the labor it involved was broken down into algorithmic steps that would ultimately be programmed into computers. Simon Schaffer (1994) traced this process through the history of the mechanization of computation, namely in Charles Babbage's attempts to build the first fully programmable calculating machines. The Babbagian project sought to mechanize the labor of computation as it had already been rationalized by Gaspard de Prony in revolutionary France where he applied Adam Smith's theory of the division of manual labor to the computation of logarithmic tables. As Loraine Daston (1994) highlights, calculation went through a "declassing" before being automated: it is because computation was turned into a mechanical and "unintelligent" human activity that machines could be expected to automate this activity. The underlying belief driving this shift was that the same *function* could be performed and executed by different sociotechnical assemblages, and that ideally the best assemblage would be the most automatic one (i.e., the most unequivocal and least error-prone one). Social activities had to be boiled down to their algorithmic steps, written into code, hardwired into mechanism, for automation to be evaluated in terms of machines outperforming humans at the same tasks.

The locus of intelligence thus shifted over time from the worker's embodied skill and the tradition of its craft to the manager overseeing the execution of an algorithmic process and the coordination of its component parts. In his reading of the Babbagian moment, Schaffer sketches out this "geography of intelligence": "Defining the *site* of intelligence was a key political task. Critics reiterated their suspicion that automatic machinery and factory discipline mechanized the proletariat. […] There was thus an unresolved contradiction between stress on the subordination of this mechanization of workers' intelligence and on the coordination and thus cerebration of their labour." (Schaffer, 1994, p. 222) Computation was gutted of its situated know-how (what we would today call embodied, enactive, or extended cognition), in a process of algorithmic rationalization so that "the very fact that a machine could execute the algorithms disqualified them as expressions of higher intelligence or rationality." (Erickson, 2013, p. 44)

This brief historical outline helps substantiate Harry Collins' qualification of late-twentieth century computers: "[t]he organism into which the intelligent computer is supposed to fit is not a human being but a much larger organism: a social group." (Collins, 1990, p. 14) The general-purpose computer, as it later evolved, is less the prosthesis or amplification of a specific organ (i.e., the brain) – as a traditional reading of technology would have it – than it is the objectivation of a certain social division of labor and intelligence.

The automation of computation reads as a trope that, by and large, holds as a general theory of technological invention and transformation, especially when it comes to forms of industrial mechanization: for a labor process to be mechanized it had to have already been formalized to resemble behaviors that could be programmed into machines. This means that mechanism, as a form of behavior and social organization, preexists mechanization by/in machines. H. Collins captures this interplay between human and machine behaviors when he points out that "there is a large aspect of human behavior that mimics machines, and machines can mimic these aspects perfectly." (Collins, 1990, p. 9) I would, however, take issue with the premise that machines perform certain behaviors *perfectly*. As French philosopher of technics, Gilbert Simondon, showed, the "perfection" we expect from machine automaticity is derived from a labor-centric evaluation of technology infused with a "yield/output morality" (Simondon, 2013). A perfect machine is akin to a compliant worker, one who flawlessly obeys commands as it executes its programmed behaviors all the while increasing productivity. "We want our robots to have all of the qualities that masters look for in their slaves, bosses in their employees, commanders in their soldiers, and we want them to have neither their weaknesses, nor their shortcomings, nor especially that irrepressible tendency towards insubordination, independence and to do as one pleases that characterizes human workers." (Dumouchel & Damiano, 2016, p. 10 [my translation]) An ideal machine does not or rather *should not* commit errors. This a problem for living organisms confronted with conflicting values and norms in action (Canguilhem, 2009; Bates ,2016). However, such a machine is ultimately an abstract ideal that has little to do with any concrete functioning or operations. Real machines break down, require maintenance, upkeep and care, real machines exist within social relations and not simply as function maximizers. The fact that machines could exhibit forms of recalcitrance or ungovernability confounds the Modern equation in which mechanism confirms the human intellect's power (who is in fact the master's power) to order reality according to the knowledge it has of its laws.

If we follow Simondon, both the industrial worker and automatic machine are subject to a functionalist labor regime whereby every individual (be it human or machine) is assigned a given position within an abstract and programmed order. It is no surprise that the master-slave analogy has so consistently irrigated the language and values of computer engineering communities (Eglash, 2007). Yield/performance morality is obsessed with the optimization of productive behaviors. Yet, as Simondon suggests, this morality does not stop at labor, but spills over into leisure and learning which become forms of production (Simondon, 2013, p. 345). In one of his more radical passages, Simondon calls upon a new form of humanism that would lift both the worker and the machine out of this yield morality (Simondon, 2013, p. 355)

3 MACHINE LEARNING AS A SOCIAL ACTIVITY

The tension between human and machine intelligence in the social division of labor is clearly palpable in the organization of contemporary labor by platforms such as Amazon's Mechanical Turk, which offers collective, distributed and extensible workforce solutions to companies with particular tasks that brute computing cannot easily or affordably manage. These Human Intelligence Tasks (HITs), as Amazon calls them, involve "Turkers" (i.e., uncontracted and outsourced laborers working from their personal computers at all hours of the day and night)[1] clicking on images to label the training data used by supervised machine learning models (e.g., is this a picture of a cat or a dog? Is this content violent or sexually offensive?). This "click-labor", which involves highly segmented and repetitive behaviors, is not so different at first glance from

---

[1] In Amazon's own terms: "As a Mechanical Turk Worker you: Can work from home. Choose your own work hours. Get paid for doing good work" and "As a Mechanical Turk Requester you: Have access to a global, on-demand, 24 x 7 workforce. Get thousands of HITs completed in minutes. Pay only when you're satisfied with the results".

the behaviors involved in the Taylorist assembly line and contrasts starkly with contemporary "creative" or "innovative" work (Irani, 2013) supposedly inherent to digital culture. Yet, while nobody would have claimed that tightening the same bolt all day could amount to a form of distinctively human intelligence, Amazon maintains that these repetitive micro-tasks exhibit a form of "human intelligence". This qualification might either be interpreted as cynicism from the world's largest and most powerful techno-logistic platform or could be taken as a provocation to ask where the intelligence of human-machine interaction resides. Much of contemporary discourse on artificial intelligence and machine learning rests upon the invisibility of the material infrastructures, situated sources of data and human labor that go into training algorithms, maintaining datacenters and curating datasets (Crawford, 2021; Sadowski, 2018)

From a socio-economical perspective, there is no doubt that the intelligence at stake in these HITs, is not that of the individual human MTurker, but the aggregated intelligence of the combined MTurkers working on a specific task. Amazon's Mechanical Turk works rather like a remote and distributed version of a factory: it derives its surplus value from the concentration and coordination of fragmented units of labor (Cardon & Casilli, 2015, p. 92). As a platform, Mechanical Turk makes otherwise unintelligent behaviors intelligent, that is it gives them a distinct exchange *value.* The extraction of surplus value by aggregating labor into intelligence rests upon industrial ideals of efficiency and optimization (Wise & Smith, 1989; Rabinbach, 1992). In other words, it is not a given activity itself that is deemed intelligent, but its coordination within an ordered process of production, an algorithm of sorts. As a mediator between theory and practice, the algorithm came to stand for Modern rationality (Totaro and Ninno 2014) and would find new expressions in the cybernetic emphasis on information and feedback as mechanisms of control and communication (Mindell, 2002; Wiener, 1961; Pickering, 2010).

The Mechanical Turk platform is, of course, a sly reference to the eponymous late 18th century automaton, which was the first "successful" mechanical chess player. It toured the royal courts of Europe as well as popular fairs, bemusing spectators, including Edgar Allan Poe and Babbage. However, the chess player's intelligence revealed the ingenuity of its mechanical design while hiding the human activity its mechanism required to function: a human chess player was in fact hidden under the board and was moving the pieces. The onlookers were willing to believe in the contraption's mechanical intelligence because they were led to focus on only part of what was happening. Similarly, Amazon's Mechanical Turk conceals the human labor involved in providing "intelligence" to its clients that it assembles and coordinates from the distributed and fragmented human behaviors. It might be said that machine intelligence – but this is perhaps true of any intelligence – is only as good as the labor it hides (Schaffer, 1996). What is labor in a "learning society" or "knowledge economy" if not learning? Paradoxically, the McLuhanian argument *qua* prophecy provides a back door for the automation of learning while claiming that automation frees up human learning to become the main form of employment.

What the history of automation shows us, is that the mechanization of labor does not unconditionally liberate humanity's "true potential" or simply free up more time for leisure. Such a "liberation" comes at a cost: *the invisibilization of the labor involved in developing automatisms*. Carnino and Marquet (2019) offer an excellent account of how the "myth of automation" covers-up the persistent failures, breakdowns and errors in contemporary technical systems such as "cloud" computing data centers, as well as the human labor involved in remedying those failures and reestablishing the external appearance of technical automaticity (i.e., functioning perfectly). The spectacle of automaticity is only as good as the stage wings it manages to conceal.

Rather than "freeing" humans from repetitive tasks, automation maintains labor as something mechanical and devoid of intelligence. It neglects the inextricably social and embodied relationship between labor, learning and intelligence as labor supposedly translates into

automatic behaviors devoid of intelligence, consciousness or reflexivity (Bates, 2016b). However, it is not the repetition of gestures and behaviors in and of itself that is unintelligent or tedious (many games rely on repetition and automaticity as mechanisms of play), but the fact that this repetition is often regulated by a dissociated and asymmetrical intelligence which, in a capitalist mode of production, is able to extract surplus value from this accumulated repetition.

To better grasp how machine behaviors interact with human ones beyond the one-sided reading of automation, a broader anthropological perspective needs to be adopted. The French paleoanthropologist of technics, André Leroi-Gourhan provided a generative distinction between three levels of human "operative chains" or behaviors: the automatic, the machinal and the lucid (or conscious). Crucially, machinal behaviors are not those that are carried out by machines, but those human behaviors which are neither purely kneejerk reactions, nor deliberate forms of attention and learning. Machinal behaviors occupy the "mental penumbra" (*pénombre psychique*) in which they disappear from consciousness throughout the normal course of events but reappear when problems or incompatibilities arise in practical situations. The ability to overcome these obstacles of daily life on an individual scale simultaneously define the locus of what is considered "intelligent" or "conscious" behavior within social practice (Leroi-Gourhan, 1965, p. 28-29).

While classical accounts of automation tend to claim machines do what is nearly devoid of intelligence in humans, machine learning shows how limited this appraisal is. Following Leroi-Gourhan's distinction, we might consider machine learning systems as social activities in which machines and humans share and develop a certain repertoire of behaviors. The point here is not to claim that automation, as a normative objective, is necessarily irrelevant – in many instances it might be a perfectly desirable, helpful, or reasonable one to pursue – but simply that it cannot be the *only* norm involved in evaluating technologies characterized by a relative open-endedness in their behaviors. Intelligence cannot be derived from automaticity. Its exercise requires behaviors that are neither completely automatic nor completely conscious. Considering machine learning developments, what Leroi-Gourhan called a mental penumbra might better be termed a sociotechnical penumbra or a "technological unconscious" (Thrift, 2004). Who is aware of "unexpected events within the sequence", who has the means of inventing solutions to problems, who gets to remember the technological unconscious and who gets to actively engage with machine behaviors and not simply react to their stimuli, are fundamental questions for tracing this emerging "geography of intelligence".

Automation presents at least three interrelated problems when applied to learning. First, and regardless of the learning theory one goes by, learning is an essentially diachronic process that involves the development of socio-technical and cognitive capacities. Automation, on the other hand supposes a compression of time whereby operations are merely the unfolding of programmed steps and time is merely the computational resource needed to execute those steps (Bachimont 2010). Second, one can never learn in someone else's place, since learning is a *relatively autonomous activity* that cannot be reduced to simple "norm-following" or executing (as opposed to "norm-establishing") behavior (Barandarian & Egbert, 2014) as is the case with automatic behavior. Third, learning's relative autonomy can only be understood in terms of its socially embedded and instituted nature – i.e., we always learn *how, to, from, with…* What and how we learn always translates what social groups value and have instituted as the ends and means of learning.

Somebody as foundational for the history and mythology of machine learning as Alan Turing was already extremely lucid about these questions when he wrote in his 1948 "Intelligent Machinery" paper about how "unfair" it would be "to expect a machine strait from the factory to compete on equal terms with a university graduate" who has been developing their intelligence over time and through social practices of learning (Turing, 1948). Turing also insisted upon the idea that if machine intelligence were going to make any sense it would have

to exhibit forms of behavior which erred and did not simply fall into automatic repetition of its programmed norm (Turing, 1950; Bates, 2016). Thus, the experience of error as positive dynamic and not simply an error to be corrected, is essential to any notion of learning.

It is precisely *because* machine learning techniques encounter failures and make mistakes that they are so powerful and enticing. It is precisely because their behaviors are partially intractable and opaque that they involve constant feedback mechanisms and couplings with engineers, designers, and users (Seaver, 2017; Grosman & Reigeluth, 2019). This tension between automaticity and intelligence is particularly palpable in contemporary ethical and social debates surrounding the opacity and biases in machine learning. They provide a clear and present example of how and why machine learning cannot be understood solely as a form of automation of human behaviors but as a social activity involving certain distributions between human and machine behaviors which do not always fit our preconceptions of how they should be distributed.

4 THE BIAS PROBLEM AND MACHINE EDUCATION

Without engaging directly in the discussion surrounding the need for algorithmic transparency (Pasquale 2015; Koliska & Diakopoulos, 2019; Sandvig, 2016) and the shortcomings or normative problems of such transparency (Annany & Crawford, 2018; Seaver, 2013; Burrell, 2016), I would simply like to consider "the elephant in the room" when thinking of algorithmic bias: the idea that learning machines just might require a form of education, if we are truly serious about them learning.

Based on this preliminary framing, we can start to look at the problem of algorithmic biases developed in supervised, reinforcement or unsupervised learning a bit differently. While biased performances in supervised and reinforcement learning should not surprise us, insofar as they are closer to forms of drilling or behaviorist conditioning, the biases developed in unsupervised learning or deep learning appear more surprising for processes where the human is supposedly more or less "out of the loop". As Michael Castelle (2020) points out in his Bourdieusian reading of Generative Adversarial Networks (GANs), there is a sense in which these types of algorithms not only implicitly or explicitly express programmed biases but may generate new ones, in the richer sense of social reproduction (i.e., social structures maintaining themselves through the variation of individual behaviors). "Put simply, GANs reproduce bias not just through their facility for stereotyped classification, but through their potential for generating *new* biased data. They differ from the "algorithms of oppression" of Google's search and recommendation engines, whose biases also exist, but which must be taken up and reproduced by humans in the loop taking practical action of their own" (Castelle, 2020). In this sense GANs manifest a degree of "technological intentionality" that emerges precisely in the open-endedness and emergent complexity of its behaviors that cannot be reduced to human intentionality through use or programming (Mykhailov & Liberati 2022).

Regardless of the specific kind of machine learning techniques we look at, we cannot avoid the perplexing fact that algorithms tend to learn *meaningful* biases, that is ones we recognize as such and that are not random or unfamiliar to us. What seems to upset us from a moral and epistemological standpoint is that algorithmic biases mirror our own biases when they are supposed to neutralize them – according to the prevalent discourse that technological processes are more rational and objective than the human ones they replace or assist. As such, bias in machine learning processes cannot be treated simply as a mathematical-logical problem[2]

---

[2] "[o]ne way to understand overfitting is by decomposing generalization error into *bias* and *variance*. Bias is a learner's tendency to learn the same wrong thing. Variance is the tendency to learn random things irrespective of the real signal." (Domingos, 2012)

that *can* have social and political ramifications when algorithmic systems go awry. Instead, it is an intrinsically social problem that can be formalized through machine learning processes. Machine learning systems are bound to develop biases (Kudina & de Boer, 2021), precisely because they are not fully automatic but relatively (depending on their degree of autonomy) open to their environments and behaviorally plastic in their responses. Just as the biases a human learner develops reveal the deeper and often cloaked stereotypes and prejudices embedded in social institutions and practices (Wellner, 2020), so too machine learning biases should draw our attention to the kinds of places, people and projects that develop them. Being able to perceive one's own biases requires a normative perspectivism (Grosman & Reigeluth, 2019) in which there are multiple optimal outcomes to a given situation. This perspectivism is impossible to foster so long as algorithms are expected to learn "on their own" and that their primary function is to optimize a given objective function established by relatively huddled engineering communities (Reigeluth & Castelle, 2021). The fact that Microsoft, for example, partially blamed ill-intentioned users for gaming its chatbot's, Tay's, "vulnerability" and inducing it to develop racist and antisemitic discourse, underlines just how limited the dominant understanding of "user interaction" is (Reigeluth, 2023).

Interestingly, the ways in which algorithms reproduce bias are usually presented as accidental or as the algorithms being "unruly" and "misbehaving", which illustrates: i) how attached we are to the image of the machine as a docile worker and how much we fear "losing control" over it (Winner, 1977), and ii) how the conjunction of learning and automaticity constantly produces *double binds* when it comes to the normative effects of algorithmic systems. Since machine learning behaviors are ever more intractable, unpredictable and analytically opaque (i.e., it is harder and harder for engineers to accurately predict how their algorithms are going to behave and to account for the actual computational steps which led to a questionable or problematic output), it might be worth approaching the engineering of these algorithms as an education process which involves not only the machines but also the engineers and designers (Winner, 1995) and users who take part in shaping algorithmic behaviors. The ways in which algorithmic learning processes are evaluated might be *reclaimed* by the plurality of agents involved in the social practices, institutions, and contexts within which they are used and upon whom they produce effects. Considering how algorithms behave in play, that is in situations where specific rules only serve as the bounds of given social interactions and not as the explanation of the behaviors themselves, it could offer a lively perspective for embracing this plurality of agents and objectives.

5 PLAYING OUT MACHINE LEARNING

Machine learning serves as a technological staging of latent yet structural social relationships that do not simply disappear because putatively more "neutral", "objective", or "rational" technologies are used. Rather, these technologies in fact reproduce, amplify, and perhaps produce new biases embedded within the very social structures where they are developed, trained, tweaked and optimized. Machine learning thus stages the conflicting norms, values and expectations that emerge around algorithms' design, implementation, and use. This brings us to the question of how error is evaluated and the status it is given within the learning process. I would like to approach this question by considering Adrian Mackenzie's distinction between "function-execution" and "function-finding" in algorithmic processes and build on that distinction by dovetailing it with the one Simondon makes between "learning" and "training".

The history and development of machine learning as a subfield of AI, which has now become nearly synonymous of AI, enfolds dimensions of play, learning and labor as normative regimes for evaluating the performance and utility of specific machine learning applications. Some of the most remarkable advances in machine learning have built upon research and applications related to human leisure and game-playing activities (e.g., checkers, chess, Go, Atari games,

Starcraft 2, etc.) (Yannakakis & Togelius, 2018) and the pre-history of AI was dramatized, as we saw, by the spectacle of automata, of which the Babbagian calculating engines and the Mechanical Turk chess player stand respectively as ideal types of labor and play. The prevalence of games in the history of AI even led figures such as Herbert Simon to submit chess as AI's "drosophila" (Chase & Simon, 1973; Ensmenger, 2012). I would like to focus here on how play, as a form of engagement with and through technology, gives us an alternate path to understanding some of the social implications connected to machine learning's large-scale deployment; one that does not presuppose an ontological rivalry between humans and machines inexorably unfolding through the cumulative process of automation (i.e. machines will replace human activities as soon as they can do those activities better than humans), but one that understands machine learning as a sociotechnical assemblage (Mackenzie, 2017; Brooks, 1990).

Consider, for example, the extent to which rule-based games have been used throughout AI's history – from Turing's imitation game to Arthur Samuel's checkers programs and Deep Blue or Alpha Go. The prevalence of computer gaming culture was one of the prime vectors for the development of applied algorithmic techniques from the 1960s onwards, wherein programming and play were synonyms for experimenting with the computer's behaviors and fostering forms of engagement with technology essential to its popularization (Triclot, 2017). On a more material level, computer gaming has contributed essential innovations to the recent renewal of connectionist machine learning, namely Nvidia's graphic cards which allowed for the parallel processing of high-dimensional matrices (Grosman & Reigeluth, 2019).

Rule-based games not only offer a sweet-spot between formal rules and creative behavior for experimenting on machine learning models (Ensmenger, 2012), but are also dramatic engineering feats that can be exhibited through spectacle and capture the public's imagination in ways that go beyond the mechanization of labor, precisely because the machines are doing something they are not supposed to be doing, or that nobody expected they would be able to do (Binder, 2021). While chess, for example, ultimately proved to be a computationally solvable problem, the fact that a computer beat the best human champion late in the 20th century does not mean humans have stopped playing chess between each other or against machines. As Ensmenger (2012) states, the fact that i) the minimax algorithm was linearly scalable with increased processing power, ii) that the game of chess had already been extensively codified (one might even say algorithmicized) by human players before ever being computerized iii) and that there were already clear benchmarks for evaluating improvement, made it look like a brute force victory over a human expert would cement its claim to a more general form of artificial intelligence. However, "other than making incremental improvements to the minimax algorithm, computer chess failed to deliver on its larger promise as a tool for exploring the underlying mechanisms of human intelligence." (Ensmenger, 2012, p. 23) As McCarthy suggests, "Computer science chess has developed much as genetics might have if the geneticists had concentrated their efforts starting in 1910 on breeding racing Drosophilia. We would have some science, but mainly we would have very fast fruit flies" (McCarthy 1997, p. 1518). Beyond these epistemological shortcomings, what the emphasis placed on sheer computational power and speed shows is how consistently the norm of automation and its yield morality guides the engineering of technology and how, in this case, this emphasis neglected the actual social activity at stake: *playing* chess and not chess as formal problem space that could be reduced to a zero-sum game in which each player is simply minimizing at every turn their maximum possible losses, as the minimax algorithm posits. In other words, testing the algorithm in an abstract, rule-bound space became a way for improving the algorithm's technical performance, but not for replacing the experience of playing chess.

In its classic definition, an algorithm protects us from error and certifies outputs based on the formal procedure that was followed (Bachimont, 2010). Conversely, the epistemic and moral

appeal of machine learning is that, through its learning, it will teach us something new about the data we feed it; that it will reveal hidden patterns and bear witness to the future as it unfolds. Machine learning algorithms are expected to *find* functions and not simply *execute* functions (Mackenzie, 2015) in ways that make them irreducible to the prior formal knowledge we may have of them and that make them performative agents embedded within social activities. Mackenzie emphasizes the social nature of machine learning, which necessarily involves forms of labor, style, and valuation. Finding a function is never merely a mathematical-logical problem, but an inherently practical one insofar as there must be some idea of how we know the function has been found and learning has thus been achieved (Mackenzie, 2015, p. 436)

The idea however, that function-finding is what confers algorithmic systems their specific learning capabilities is somewhat at odds with the overwhelming emphasis that is placed on optimization as a measure of algorithmic performance within machine learning communities and practices. While learning and optimization seem rather synonymous at first glance – indeed both denote an ameliorative performance of sorts – there is a sense in which they are barely synonyms at all. As a rule, machine learning algorithms are evaluated based on their ability to minimize a loss function or maximize an objective function, which serves as the mathematical-behavioral norm to be achieved through learning. Although it is not clear in advance *what* it is they will learn (i.e., the specific correlation or prediction), formal criteria exist to determine whether the algorithm has improved its performance (i.e., it has achieved a better approximation of the objective function).

6 TRAINING VERSUS LEARNING

While evaluating learning based on standardized metrics may satisfy a minimalistic definition of learning, improved performance on a test certainly does not meet a socially and subjectively meaningful definition of learning in which the test is, at best, a way of determining what still needs to be learned, but not what has been learned and what that learning means for the learner (Vygotsky, 1987). Evaluating machine learning based on its optimization of an objective function is basically equivalent to the emphasis placed on skill acquisition measured by standardized testing. Just as the labor of calculation was broken down into a social algorithm before being mechanized, so too learning was already practically turned into an optimization problem by the widespread use of standardized testing and certification in human education before it could be presented as such by the machine learning community.

The discrepancy between optimization and learning can be further problematized by way of a useful conceptual distinction made by Gilbert Simondon between two different modes of adaptation: training (*dressage*) and learning (*apprentissage*). Simondon, who was well-versed in cybernetics and information theory, developed this distinction throughout the 1960s and into the 1980s, when debates around the ontological and epistemological limits to artificial intelligence and learning were intense and transdisciplinary. Simondon recognized that computers exhibited complex forms of adaptive behavior as they processed information but was skeptical that this performance could be equated to a form of learning, or that this learning was something that could be automated. One of Simondon's main contentions is that the feedback processes that characterize "open machines" whose behaviors are not entirely pre-programmed do not eject humans from the loop, but in fact provide a new function for human regulation and intelligence. Training is characterized by a "convergent determinism" and learning by a "divergent determinism" (Simondon, 2013, p. 345). The former limits the impact new information has on a system by maintaining a certain functional stability (e.g., typically any negative feedback mechanism that ensures incoming information does not modify the system's equilibrium beyond preset bounds). Conversely, learning overrides normative bounds by continuously opening itself to new forms and sources of information. Training involves

behavioral adaptation to meet an expected output or goal while learning is constantly reinventing the activity's very goal. Simondon makes this epistemological distinction with a more social and political one in mind between the machine and the slave. The slave can hypothetically always reinvent the purpose of its action, while even the most sophisticated computational machine can only adapt to incoming information (Simondon, 2013, p. 344-346). Simondon stresses that the human is a rather "dangerous automaton" because it continuously threatens to override established norms, which of course is precisely the opposite of what we expect of a perfect automaton to do, that is execute a norm or function without erring, protesting, or changing course. In light of this distinction, it might then be more precise to qualify most current machine "learning" applications as machine *training*.

Building upon a cybernetic approach to learning, the political scientist Karl Deutsch made a similar distinction as Simondon's between "learning" in which feedback mechanisms simply converge towards an objective and learning through which the objective itself is displaced in the process (Deustch, 1966, p. 92). Interestingly, Deutsch uses this distinction to refute the dominant game-theoretic premises that games are played based on a shared knowledge of the rules by the players and that the rules of the game cannot be changed by the action taken. "Just as game theory ordinarily does not allow for changes in the performance characteristics of particular elements, so it does not provide ordinarily for changes in the rules of the game. Taken together, these two restrictions seem to cut it off from the description of much of the process of learning." (Deutsch, 1966, p. 57) Castelle's more recent critique of the game-theoretic underpinnings to GANs reminds us that while the connexionnist spirit of contemporary AI research is generally presented as a refutation of symbolist and rule-based cognitivism, "the dependency of neural network training on loss/utility function means that deep learning and GANs maintains a very close link with instrumental rationalism (i.e., players follow the rules to maximize their preferred outcome) (Castelle, 2017). In other words, there is still a penchant for mistaking the rules of the game for the game itself understood as an iterative, generative, and always partially unpredictable social activity.

7 WAR GAMES

The idea that human learning is constantly transforming the very objective of its learning was something Norbert Wiener clearly outlined when he wrote about the necessity of programming a learning machine for the simulation of nuclear war games to avoid the catastrophic outcome of a nuclear war (Wiener, 1961, p. 177). Victory here serves as the objective function to be maximized, but to what end? Can a war ever be properly simulated as a zero-sum game? We could formalize all the conditions whereby victory would be granted, but declaring victory is only ever as good as the social relationship whereby victory is conceded by the other party and *means* victory to those playing, be it in war, elections, kick-the-can or chess. We are faced with the inescapable and insolvable need to account for what we mean by "victory". Simply following games' rules does not entail the game will be won. There is always a certain social relationship involved. Even when one plays by oneself as is apparently the case with single player computer games: the player is in fact engaging with the game's AI which embodies the social relationships and behaviors the game is supposed to simulate (Ash, 2015; Triclot, 2017). Moreover, while the computational simulation of war has consistently infused the history computer games, from *Spacewar* to *Call of Duty*, nobody would contend that such experiences of play are comparable to the experience of war itself. This points to the paradoxical and conditional nature of play as Bateson famously showed in his analysis of wolf cubs biting playfully: "These actions in which we now engage do not denote what those actions for which they stand would denote."(Bateson, 1987, p. 139) Nibling is not a preparation to bite to wound or kill, it is something else entirely and yet it can become that at any moment, should the conditions of play fall apart. Play appears as both a form of social activity that breaks with the normal course

of events, as a distinct rule-bound space (Huizinga, 1951) and as the "self-obligating attitude" (Henriot, 1969) or form of "experience" (Caillois, 1967) latent within everyday activities and which brings the rules of the game to life. In other words, playful engagement is not simply about executing rules coded into an algorithm but about generating new forms of behavior that bend the rules into practice. Those familiar with Ludwig Wittgenstein's concept of "language-games" or Pierre Bourdieu's theory of habitus will find heuristic connections in this framing of game-playing. The French word for game (*jeu*) denotes all the combinable elements of a game (e.g., "*un jeu de cartes*") as well the spaces of indeterminacy in a mechanism (e.g., "*il y a du jeu dans le mécanisme*"), as in a doorhandle that needs to be jiggled to be opened (something we might call "drift" or "leeway"). This semantic ambivalence is helpful to characterize processes such as machine learning in which there is a space between the formal rules and their effective outputs, a space filled by playing out the possibilities, by developing styles and ways of doing things, by building repertoires of machinal behaviors. While this drift is not new, who defines its threshold for correction and how it is corrected, in part defines the locus of intelligence in these sociotechnical assemblages.

8 DIVERGENT AND CONVERGENT PLAY

The behavioral complexity of play thus involves both the capacity to follow rules (be they explicit or tacit) and to improvise upon their combination and potentially invent new moves or styles of play. When compared to the normative regime of automation, play tells us at least two important things in terms of the expectations placed on machine learning.

First, nobody would expect computers to completely replace human-players in a game of chess or checkers: the whole point is to see what kind of specific challenge and pleasure can be derived from playing against a machine or watching a machine and a human play. The immediate pleasure and function of play is playing (Huizinga, 1951; Bateson, 1987; Massumi, 2014), which cannot be said of automation where the value and utility lies outside the activity of labor itself. While even the most "low-skilled" jobs (e.g., cashiers, waiters in fast-food restaurants, cleaning jobs, etc.) have proven extremely complex to automate (Collins & Kusch, 1998), there tends to be a general, if implicit and rather murky, consensus around the idea that these jobs *should* ultimately be automated so as to free up workers for learning and leisure, like McLuhan had prophesized. This is intriguing considering contemporary accusations leveled upon social media, digital platforms, and nudging technologies of "infantilizing" users by "gamifying" their engagement and fostering forms of unintelligent, regressive, or addictive behavior. The core of this charge lies in idea that the gamification of user engagement covers up the process of surplus-value extraction of users' attention – hence the idea that users are in fact working by simply paying attention to certain applications. While gamification in and of itself is perhaps not to blame for this form of relationship to digital media, it is certainly ironic to see millions of users during their commutes, at home or in checkout lines, interacting with their smartphones with gestures that are about as limited and repetitive as industrial factory workers', all optimizing the platform's objective function.

Second, and in contrast with this contemporary preoccupation around "gamification", play, as form of social activity, points to a much different relationship to AI and, more generally, to technology. Whereas organized labor predicates technology's value on its ability to maximize output functions by optimizing production processes, play involves a more open-ended engagement and use of technology, in which the machine is not simply seen as a tool, a means to an end, but as a partner in action through which new behaviors can be invented, often seamlessly or unwittingly (Winnicott, 2005). We might consider the human-machine interaction involved in recommendation algorithms used by platforms such as Netflix or Spotify as a paradigmatic case of a ludic interaction. The "collaborative filtering" techniques seek a sweet spot between items users have already consumed and ones they are likely to enjoy based on

what similar profiles have consumed. Users are telling algorithms what they want; algorithms are predicting what users want; and both are playing out an iterative experience of culture (Reigeluth, 2017) akin to a certain form of mimesis evocative of Leroi-Gouran's machinal behaviors. Interestingly, this is also perhaps one of the most publicized examples of platforms keeping their algorithms secret to foster "natural" interactions where users are not trying to "game" the system (Cardon, 2013) – *A/B testing* deployed on digital platforms generally operates according to the same normative camouflage when trying to measure the uptake by users of two different version or functionalities.

Again, the normative stake that keeps rearing its head revolves around the fact that it is those who can design and predispose these different behavioral interactions to maximize an objective function who have a claim to the system's intelligence. The drift or divergence inherent to play is constantly bounded, corrected, realigned by the rigidness of the rules and the strength of social relations involved. The issue, however, is not the boundedness of play in and of itself, but the power relations and types of collective behavior they reproduce and reinforce. Play is a disruptive yet necessary form of engagement for the logics of any system seeking to "trap" users (Seaver, 2019) and yet, play itself seems to challenge the very efficacy of recommender systems (e.g. the Amazon recommender algorithm cannot differentiate between "serious" and "playful" queries; YouTube cannot differentiate if you watch certain politically themed videos mockingly or approvingly). The continual improvement of recommender algorithms does not simply rely upon our automatic responses to pairwise stimuli, nor the exhaustive formalization of all possible behaviors in the AI. When we interact with machine learning systems in the "normal" sequence of behaviors, we do not distinguish our behaviors from the systems', we are part of the system, folded into its "envelop" constantly soliciting our attention and engagement (Ash, 2015). Algorithms might be "black boxes" for users, but so are users for algorithms; it is the interplay of their partially indeterminate behaviors that tells us if something is being learned or if their behaviors are merely converging towards a preestablished norm. By taking up Simondon's distinction between training (convergent determinism) and learning (divergent determinism), we might then distinguish between two kinds of play: divergent and convergent play. While the former might be seen as unproductive, wasteful, useless when looked at from the labor-centric automation paradigm, it helps us understand why complex sociotechnical assemblages such as machine learning have a nasty habit of drifting: rules do not govern social practice as a classical algorithm does a computer. It is by playing with rules that they become effective. And machine learning systems are closer to a social practice than they are to a classical algorithm. Convergent play, on the other hand, highlights when and how this inevitable tendency to drift is realigned on an objective function. Such convergence is inevitable in social practice: there is always a moment in which dispersion is unified. If, however, game-playing has become the predominant form of engagement with contemporary technology (Flusser, 1993) and if one of the essential characteristics of play is that the players are free to enter and leave as they please (Henriot, 1969, Caillois, 1967), then we must consider the leeway these algorithmic "traps" allowing in playing with their rules and whether we are able to leave them if we cannot.

Recent debates around Open AI's chat GPT underline the normative tensions at play within such systems. Social media and academic commentary have abounded around the playful, intriguing, and sometimes troubling interactions with the system. The prompts users (if this term is still appropriate) feed the algorithm are not simple commands that execute programmed behavior, they serve as new occasions for it to enrich its rule-based model of linguistic interaction. GPT can imagine a drunken rant by Michel Foucault but can have trouble figuring out whether 10kg of feathers weighs more or less than 10kg of iron. AI scholars like Gary Marcus have been quick to point out that "these systems can be incredibly fun to play with", but how "unreliable and potentially dangerous" they could be (Marcus, 2022). It's hard not to think of Simondon's qualification of human's as rather dangerous automata and wonder if these algorithmic systems

that imitate and iterate the human behaviors from which they learn are not rather dangerous automata; and whether the rule-based but playfully open-ended engagement with them is not in large part responsible for this uncanny feeling.

9 CONCLUSION

Interacting with machine learning algorithms and framing their behavior in terms of game-playing, gives us a new dimension for thinking about what machines do when they are not simply obeying orders. By exploring some of the historical tropes behind contemporary machine learning, I showed that automation is not the only way to understand what it means for a machine to learn. Without its playful dimension, automation misses the potential inventiveness of machine behaviors within a social practice. If automation, as we have shown, perpetuates the moral economy of docility, efficiency, and control, play invites different moral criteria for evaluating learning, not merely as a means for maximizing an objective function, but recognizing the plurality of values at play. We might then ask: was the learner's solution clever, creative, or boring? Did its learning conjure beauty, pleasure, or excitement? Were the effects its learning produced upon the social activity within which it unfolded redundant, sensible, or challenging? These questions may seem far removed from many of the technical stakes machine learning involves, but I am only asking what we might ask of machines should we want them to learn, which in turn begs the question of what it means for us to learn in a time when machines are embedded in our social practices, not only as things that operate in our stead, but as things which operate with and through us.

Data Access Statement
No new data generated or analyzed.

Contributor Statement
The author is the single author of the manuscript.

References

Ananny, M., & Crawford, K (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society.* 20(3):973-989.

Ash, J. (2015). *The Interface Envelope. Gaming, Technology, Power.* New York and London: Bloomsbury Press.

Bachimont, B. (2010). *Le sens de la technique: le numérique et le calcul.* Paris : Editions Les Belles Lettres.

Barandiaran X., & Egbert, M. (2014). "Norm-establishing and norm-following in autonomous agency." Artificial Life. 20(1):5-28.

Bates, D. (2016). Automaticity, Plasticity and the Deviant Origins of Artificial Intelligence. In Bates, D., Bassiri, N. (eds) *Plasticity and Pathology: On the Formation of the Neural Subject.* New York: University of Fordham Press.

Bates D. (2016b). Insight in the Age of Automation. In: Chaplin J.E., McMahon D.M. (eds) *Genealogies of Genius*. New York: Palgrave Macmillan.

Bateson, G. (1987). A Theory of Play and Fantasy. in *Steps to an Ecology of Mind*. Nortvale: Jason Aaronson Inc.

Binder, W. (2021). Alpha Go's Deep Play: Technological Breakthrough as Social Drama. In *The Cultural Life of* Roberge, J. and Castelle, M. (Eds), *Machine Learning.* Palgrave-Macmillan.

Brooks, R. (1990). Elephants Don't Play Chess. *Robotics and Autonomous Systems* 6(1,2), 3–15

Burrell, Jenna. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*. doi:10.1177/2053951715622512

Caillois, R. (1967). *Les jeux et les hommes.* Paris : Gallimard.

Canguilhem, G. (2009). Machine et organisme. In *La connaissance de la vie.* Paris : Libraire philosophique Vrin.

Cardon, D. (2013). Inside the Mind of PageRank: A study of Google's algorithm. *Réseaux*, 177(1), 63-95. https://doi.org/10.3917/res.177.0063

Cardon, D, & Casilli, A. (2015). *Qu'est-ce que le Digital Labor.* Bry-sur-Marne : INA Editions.

Carnino, G., & Marquet, C. (2019). Du mythe de l'automatisation au savoir-faire des petites mains : une histoire des datacenters par la panne. *Artefact,* 11, 163-190.

Castelle, M. (2020). The Social Lives of Generative Adversarial Networks. *ACM Conference on Fairness, Accountability, and Transparency.*

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4(1), 55–81.

Citton, Y. (2017). *The Ecology of Attention.* Cambridge: Polity Press.

Collins, H., & Kusch, M. (1998). *The Shape of Actions : What Machines and Humans Can Do.* Cambridge, MA: MIT Press.

Collins, H. (1990). *Artificial Experts: Social Knowledge and Intelligent Machines.* Cambridge MA: MIT Press.

Crary, J. (2013). *24/7: Late Capitalism and the Ends of Sleep*. London and New York: Verso.

Crawford, K. (2021). *Atlas of AI.* New Haven: Yale University Press.

Daston, L. (1994). Enlightenment Calculations. *Critical Inquiry,* 21(1), 182-202.

Deutsch, K.W. (1966). *The Nerves of Government, Models of Political Communication and Control.* New York: The Free Press.

Domingos, P. (2012). A Few Useful Things to Know about Machine Learning. *Communications of the ACM*. https://doi.org/10.1145/2347736.2347755

Domingos, P. (2015). *The Master Algorithm*: *How the Quest for the Ultimate Learning Machine Will Remake Our World.* Allen Lane, *SineLoco.*

Dumouchel, P., & Damiano, L. (2016). *Vivre avec les robots.* Paris : Editions du Seuil.

Eglash, R. (200)7. Broken Metaphor: The Master-Slave Analogy in Technical Literature. *Technology and Culture,* 48(2), 360-369

Ensmenger, N. (2012). Is chess the drosophila of artificial intelligence? A social history of an algorithm. *Social Studies of Science*, 42(1), 5–30

Erickson, P., Klein, J.L., Daston, L., Lemov, R., Sturm, T., & Gordin, M.D. (2013). *How Reason almost lost its Mind.* Chicago: University of Chicago Press.

Flusser, V. (1993). *Pour une philosophie de la photographie.* Circé Press.

Friedmann, G. (1946). *Problèmes humains du machinisme industriel.* Paris : Gallimard.

Grier, D.A. (2007). *When Computers were Human,* Princeton: Princeton University Press.

Grosman, J., & Reigeluth, T. (2019). Perspectives on algorithmic normativities: engineers, objects, activities. *Big Data & Society*. doi:10.1177/2053951719858742

Henriot, J. (1969). *Le jeu.* Paris : Presses Universitaires de France.

Huizinga, J. (1951). *Homo ludens.* Paris: Gallimard.

Irani, L. (2015). The cultural work of microwork. *New Media & Society*, 17(5), 720-739. doi:10.1177/1461444813511926

Koliska, M., & N. Diakopoulos, N. (2019). Disclose, Decode and Demystify: An Empirical Guide to Algorithmic Transparency. In *The Routledge Handbook of Developments in Digital Journalism Studies*. Eldridge S. and Franklin, B. (eds). New York: Routledge.

Kudina, O. & de Boer, B. (2021). Co-designing diagnosis: Towards a responsible integration of Machine Learning decision-support systems in medical diagnostics. *Journal of Evaluation in Clinical Practice*, 27(3), 529–536. https://doi.org/10.1111/JEP.13535

Leroi-Gourhan, A. (1965). *Le geste et la parole. Tome II.* Paris : Albin Michel.

Light, J. (1999). When Computers were Women. *Technology and Culture*, 40(3), 455-483.

Lowrie, I. (2017). Algorithmic rationality: Epistemology and efficiency in the data sciences. *Big Data & Society.* https://doi.org/10.1177/2053951717700925

Mackenzie, A. (2015). The production of prediction: What does machine learning want? *European Journal of Cultural Studies*, 18, 429-445.

Mackenzie, A. (2017). *Machine Learners. Archeology of a Data Practice.* Cambridge, MA: MIT Press.

Marcus, G. (2022). AI's Jurassic Park moment, *https://garymarcus.substack.com/p/ais-jurassic-park-moment.*

Massumi, B. (2014). *What Animals Teach Us about Politics.* London and Durham: Duke University Press.

McCarthy, J. (1997). AI as Sport. *Science,* 276(5318):1518-1519.

McLuhan, M. (1964). *Understanding Media. The Extensions of Man*. New York: Signet Books.

Mindell, D. (2002). *Between Humans and Machines.* Baltimore: The Johns Hopkins University Press.

Moulier-Boutang, Y. (2012). *Cognitive Capitalism*. Cambridge : Polity Press.

Mykhailov, D., & Liberati, N. (2022). A Study of Technological Intentionality in C++ and Generative Adversarial Model: Phenomenological and Postphenomenological Perspectives. *Foundations of Science,* 2022, 1–17.

Noble, D. (1984). *Forces of Production; A Social History of Industrial Automation,* New York: Knopf.

Pasquale, F. (2015). *The Black Box Society,* Cambridge, MA: Harvard University Press.

Pickering, A. (1993). *The Mangle of Practice.* Chicago: The University of Chicago Press.

Pickering, A. (2010). *The Cybernetic Brain.* Chicago: The University of Chicago Press.

Rabinbach, A. (1992). *The Human Motor. Energy, Fatigue, and the Origins of Modernity*. Oakland: University of California Press

Reigeluth, T. (2023). Machine Learning Normativity as Performativity. in *Handbook of Critical Studies of Artificial Intelligence*, Simond Lindgren (ed.). Edward Elgar Publishers.

Reigeluth, T., & Castelle, M. (2021). What Kind of Learning is Machine Learning? in *The Cultural Life of Machine Learning,* Roberge, J. and Castelle, M. (eds). Palgrave-Macmillan.

Reigeluth, T. (2017). Recommender Systems as Techniques of the Self ? *Le foucaldien,* 3(1), 1-25.

Sadowski, J. (2018). Potemkin AI. *Real Life,* August 6.

Sandvig, C. (2016). When the Algorithm Itself Is a Racist: Diagnosing Ethical Harm in the Basic Components of Software. *International Journal of Communication,* 10, 4972-4990.

Sawday J. (2007). *Engines of the Imagination. Renaissance culture and the rise of the machine.* New York: Routledge.

Schaffer, S. (1996). Babbage's Dancer and the Impresarios of Mechanism. In Spufford, F. and Uglow, J. (eds) *Cultural Babbage: Technology, Time and Invention*. London and Boston: Faber and Faber.

Schaffer, S. (1994). Babbage's Intelligence: Calculating Engines and the Factory System. *Critical Inquiry*, 21(1), 203-227.

Seaver, N. (2019). Captivating algorithms: Recommender systems as traps. *Journal of Material Culture.* 24(4):421-436.

Seaver, N. (2013). Knowing algorithms. *Media in Transition,* 8.

Seaver, N. (2017). Algorithms as culture: Some tactics for the ethnography of algorithmic systems. *Big Data & Society*. https://doi.org/10.1177/2053951717738104

Simondon, G. (2013). *L'individuation à la lumière des notions de forme et d'information.* Grenoble : Editions Jérôme Million.

Steiner, C. (2012). *Automate This.* New York : Penguin Books.

Stiegler, B. (2015). *La société automatique, 1. L'Avenir du travail.* Paris : Fayard.

Totaro, P. & Ninno, D. (2014). The Concept of Algorithm as an Interpretative Key of Modern Rationality. *Theory, Culture & Society.* 31(4), 29-49.

Triclot, M. (2017). *Philosophie des jeux video*. Paris : Editions de la Découverte.

Turing, A. (1950). Computing Machinery and Intelligence. *Mind*, New Series, 59(236), 433-460.

Turing, A. (1948). Intelligent Machinery. In B J Copeland (ed.), *The Essential Turing* (Oxford, 2004; online ed, Oxford Academic, 12 Nov. 2020), https://doi.org/10.1093/oso/9780198250791.003.0016

Vygotsky, Lev S. 1987. Thinking and Speech. In *The Collected Works of LS. Vygotsky. Vol. I: Problems of General Psychology (Trans. N. Minnick)*, edited by R.W. Rieber and A.S. Carton, 39–285. New York: Plenum.

Wellner, G. (2020). When AI Is Gender-biased. *HUMANA.MENTE Journal of Philosophical Studies*, 13(37), 127–150. https://www.humanamente.eu/index.php/HM/article/view/307

Wiener, N. (1961). *Cybernetics or, Control and Communication in the Animal and the Machine.* Cambridge, MA: MIT Press.

Winner, L. (1977). *Autonomous Technology. Technics-out-of-control as a Theme in Political Thought*, Cambridge, MA: The MIT Press.

Winner, L. (1995). Citizen Virtues in a Technological Order. In A. Feenberg & A. Hannay (Eds.), *Technology and the Politics of Kowledge*). 65-84. Indianapolis: Indiana University Press.

Winnicott, D.W. (2005). *Playing and Reality.* London: Routledge.

Wise, N. & Smith, C. (1989). Work and Waste: Political Economy and Natural Philosophy in Nineteenth Century Britain (I). *History of Science.* 27(3), 263-301.

Wise, N. 2007. The Gender of Automata in Victorian Britain. In J. Riskin (Eds.), *Genesis Redux: Essays in the History and Philosophy of Artificial*, 163-195. Chicago: University of Chicago Press.

Yannakakis, G.N. & Togelius, J. (2018). *Artificial Intelligence and Games.* Springer.

Zuboff, S. (1988). *In the Age of the Smart Machine: The Future of Work and Power*. New York: Basic Books.