



Article scientifique

Article

2024

Published version

Open Access

This is the published version of the publication, made available in accordance with the publisher's policy.

Bruno Latour and Artificial Intelligence

Venturini, Tommaso

How to cite

VENTURINI, Tommaso. Bruno Latour and Artificial Intelligence. In: Tecnoscienza, 2024, vol. 14, n° 2, p. 101–114. doi: 10.6092/issn.2038-3460/18359

This publication URL: <https://archive-ouverte.unige.ch/unige:175094>

Publication DOI: [10.6092/issn.2038-3460/18359](https://doi.org/10.6092/issn.2038-3460/18359)

© The author(s). This work is licensed under a Creative Commons Attribution (CC BY 4.0)

<https://creativecommons.org/licenses/by/4.0>

Bruno Latour and Artificial Intelligence

Tommaso Venturini 

University of Geneva

Corresponding author

Tommaso Venturini
University of Geneva, Medialab
& CNRS, Center for Internet and
Society
Boulevard du Pont-d'Arve 40,
1205 Genève, Suisse
[✉ tomm.venturini@gmail.com](mailto:tomm.venturini@gmail.com)

Submitted: July 27, 2023

Accepted: October 31, 2023

Abstract

This scenario discusses generative AI in light of Bruno Latour's sociology of technology. It considers why Latour showed little interest in the simulation of intelligence and how connectionist AI fails to meet his condition for scientificity but offers a fascinating writing mediation. AI is most interesting not because it emulates human thinking or writing, but because it differs from them. Drawing on actor-network theory, this scenario argues against the idea of machines becoming detached from their creators and highlights how AIs can only exist through the support of their human assistants. The risks associated with these technologies do not come from an improbable singularity, but from their embedding in the dull and exploitative industry of digital attention economy.

Keywords

artificial intelligence; actor-network theory; machine learning; deep learning; generative AI.

1. Introduction

Since his disappearance in October 2022, I have been missing Bruno Latour intensely, as have all those who knew him. What we miss the most, I believe, is his unwavering intellectual compass, his ability to see through the “foam of events” (as Lippmann put it) and identify what matters in the long run. My ache for his compass has only increased in the last few months, as generative artificial intelligence has suddenly become *the* matter of concern in digital societies. The recent leap forward in AI's ability to perform advanced tasks (translating and writing texts, designing images, creating videos, etc.) has struck a deep chord in our collective imagination and has become the pivot of a surge of hopes and worries. Some see it as the miracle solution that will unlock unlimited intellectual progress, optimize the allocation of resources and become the faithful companion of our thinking. Others (and most social scientists tend to fall into this category) are increasingly wary of the prejudices linked to this technology, its possible abuses, its consequences on labor, copyright, public debate, education, etc. STS scholars have already begun to investigate these imaginaries and their roots (cf., for example, Suchman 2008; Katzenbach 2021; Bareis and Katzenbach 2022; Ferrari and McKelvey 2023; Veale et al. 2023).

But what would Bruno make of it? I keep asking myself: what would he think is the story here? What should we hope or fear, or dismiss as nonsense? I am sure that Bruno would know. After all, no one has so skillfully woven anthropology, philosophy, sociology, and political science to unfold the imbroglis of contemporary technology.

Sadly, we can't ask for his help on this one. All we can do is engage in a thought experiment that (somewhat ironically) resembles a peculiar ability of generative AI: we can guess what Latour might have written about these technologies. We could even ask the AI to do so. And ask I did, but (somewhat reassuringly) obtained no useful insight – though the machine did a good job of reproducing Latour's humor, though maybe in a bit of a caricature:

Ah, the generative artificial intelligence, a marvel of human creation, yet an entity in its own right, weaving its digital tendrils into the fabric of our society. As with many things, it is neither entirely of our making nor entirely separate from us, inextricably linked, yet perpetually distant... We teach them our ways, our languages, our desires, only to be confronted with their peculiar modes of thought, their uncanny reflections of our own selves. Their very existence forces us to renegotiate our place within the entangled web of actors, human and non-human alike. And thus, we find ourselves caught in a dance with these digital phantoms, a perpetual *pas de deux* between master and servant, creator and creation. And who is to say, when the music fades and the stage lights dim, which role we shall assume? (ChatGTP4 impersonating Latour)

Not bad as the style goes, but quite shallow in meaning – as one would expect from a machine that (as advertised) is good at writing by so-so at reasoning. So, in this scenario, I will make the attempt myself, scraping ideas from the few writings where Latour engaged with similar questions. To be sure, this will entail some betrayal, not only because (as Latour argued multiple times) translation is always a form of treachery (*traduction trahison*), but also because Latour did not seem particularly interested in AI in the first place.

Sifting through the vast collection of his writings, I found only one paper dedicated to AI – a theatrical dialogue with novelist Richard Powers about HAL and the Turing test (Latour and Powers 1998). He dedicated the last part of his career to the enormous threats posed by the ecological crisis and the enormous political work necessary to rebuild a common world. From this perspective, I believe he would have considered AI as a distraction, or worse he would have been appalled by the enormous quantity of energy consumed by these systems. Not unlike the dream of colonizing Mars or escaping to the metaverse, criticized in his Critical Zone exhibition, he might have regarded the efforts invested in creating future artificial intelligences as a means of procrastinating the much more difficult and urgent task of cultivating the present social intelligence.

This raises a moral (or rather political) dilemma: if Latour might have seen AI as a distraction from the pressing issue of Gaia, is it fair to invoke his thought to write about it? I believe so, as I believe that this is not just an exercise in style. For reasons that I hope will become clear, thinking with Latour (or “as Latour”) can help unfold the peculiar nature of generative AI.

2. AI's possible role in science and social sciences

Latour's relative indifference towards AI is not surprising. Throughout his career, he consistently admired science and technology for their ability to make us sensitive to natural and social phenomena, rather than for their capacity to artificially reproduce them. Latour cherished artificiality, to be sure ("the more constructed the more real", Latour 2003a), but for its capacity to renew rather than to reproduce. He was an "amateur of science" (Latour 1993a), but I don't think AI, and deep learning in particular, would have qualified as such in his mind. Famously, Latour described the scientific method as a reversible chain of transformations: "an unbroken series of well-nested elements, each of which plays the role of sign for the previous one and of the thing for the succeeding one" (Latour 1995, 56). The explicitness and reversibility of this chain of "reference" was to him nothing less than "mode of existence" of science (Latour 2013). By such a standard, I doubt that the computational blender of generative AI can be considered as scientific. Indeed, the success of connectionist AI and its leap forward from symbolic AI stems precisely from the choice of sacrificing transparency in exchange of efficacy (Cardon et al. 2018). At its best, generative AI represents the triumph of what Latour (2013) calls "double click": the ultimate success of technological systems that become invisible and inscrutable because of their own efficiency. But, at its worst, it also resembles the "double click" ideology, the erroneous belief that knowledge can be obtained immediately and at no cost.

Latour was a fan of digital methods, but remained skeptical of simulations and models (Venturini et al. 2015). To him, "The Whole is Always Smaller Than Its Parts" (Latour et al. 2012), and by "smaller" he meant "less interesting". He never wanted to compute or aggregate masses of data points, but was captivated by the idea of observing them individually. For him, this was the real advantage of digital research (Latour 2007) and the purpose of the quali-quantitative approach he proposed (Venturini et al. 2017; Venturini 2024). He began his career as an ethnographer and never gave up his penchant for keen observation and painstaking description. He never cared about distant reading and I doubt that AI distant writing would have held much appeal for him *as an analytic tool*.

On the other hand, he would have been intrigued by AI as a *writing* tool. As well-known, Latour was passionate about the practice of sociological writing, which he saw as the equivalent of natural sciences laboratories. To him, writing was our chain of reference:

If the word "scientific" can be applied to social science, and I think it does, it is above all for the writing work that we are obliged to do for each field differently. This is what makes "objective" our theses or our reports or our books. (Latour 2014, *my translation*).

Interestingly, in such a struggle, thinking and writing are inseparable, forming a practice that Latour describes with words that closely reminds the functioning of generative AIs:

You are thinking because you are writing, and the writing makes you think... Writing makes one think, because you write stupid things. On the next page, something happens, and on the next. It comes out of this weird thing [laughs] that I have a lot of trouble teaching my students in my writing workshops: "If you don't write you aren't thinking". (Latour in an interview by Coccia 2021).

Large language models indeed work by reading a prompt and then predicting the bit of text most likely to follow, and then the next and then the next (“on the next page, something happens, and on the next”).

Because of this functioning, some people argue that AIs are not really reasoning but merely writing. I think Latour would have found this distinction specious. Instead, what he believed to be unique to human writing is the capacity to return to the same text over and over again in order to reread and rewrite it, especially when this is done collectively in an in-person writing workshop or in a remote dialogue with other authors. But Latour would not have balked at the notion of introducing a machine into the mix. Quite the contrary, I believe he would have liked the idea of an interactive writing tool that is less obedient than a typewriter, as this may add another layer of resistance, and thus reflexivity, to the thinking of the researcher who uses it.

3. AI and the complementarity between human and technological agents

Another reason why I am convinced that Latour’s reading of AI renaissance would have been quite down-to-earth is that he explicitly said so, in the only text he dedicated to this technology:

I don’t understand the fuss about futuristic machines at all. Imagine a Turing test for shoveling, here is a human worker shoveling earth and here is an earth-moving machine, can you tell them apart? Yes, of course, the second one is immensely more powerful and drives tons away instead of kilos. What would be the lesson to draw? Why would we build machines if it were not to make them immensely more powerful than ourselves? To make us collectively more powerful than we were in the past. What is true for shoveling, is true for intelligence and also for emotion. Is not Hollywood such an immensely powerful machine, to be able to produce tears, passion, love and fright? How bizarre is it, to imagine a Turing test where the two halves are on an equal footing, when all what we strive for, on the contrary, is *inequality*? But why do we conclude from this inequality that machines are escaping us, and dominating us? (Latour and Powers 1998, 186, *emphasis added*).

The quote operates a classic Latourian move, challenging the exceptionality of one technology (AI) by showing its fundamental resemblance to other technologies (the shoveling machine) and societal apparatus (Hollywood as an emotion machine). It says: artificial intelligence is just another technology; it is different from previous ones (for every technique is unique in some way), but not extraordinary.

In the discussion from which this quote is extracted, the sci-fi novelist Richard Powers repeatedly tries to convince Latour of the disruptiveness of an AI capable of beating the Turing test and the doubts it casts on our own thinking. If a mindless machine can *pretend* to think, how can we be sure that humans are not doing the same?

Am I, myself, capable of passing the Turing Test honestly, or do I do so only by cheating? Is consciousness itself a form of “cheating”, whereby I grace with the most charitable interpretation the semblance of intelligence burbling up from my lower-level machines? (*ibid.*, 182).

Latour seems utterly and unsurprisingly unimpressed by this argument. After all, half of his actor-network theory (ANT) revolves around the idea that anything that acts socially is a social actor – a title that should thus rightfully be granted to machines and to all sorts of non-human beings (Latour 2005). Latour famously wrote that the actions of a speed bump, or a “sleeping policeman” as sometimes called, are not fundamentally different from that of a human policeman: they both slow the traffic down (Latour 1994a). True, the sleeping policeman has no soul, intentions or emotions, but this does not prevent him from acting in society and thus to be a social actor.

To be sure, Latour never said that human actors and non-human actors are the same. Of course, they are not. A sleeping policeman is much cheaper to install and maintain and can work 24 hours a day without complaint. That’s the whole interest of technologies: because they are different from humans, they can deal with the same tasks in ways that are different and sometimes more efficient. Precisely because their inner workings are different from those of humans, the shoveling machine can move more earth than a human shoveler and the Hollywood machine more emotions than a human storyteller. In the exact same way, generative AIs do not think, write, or draw as humans do and that is why they are interesting.

Generative AIs are most remarkable in that they act differently from humans. Like shoveling machines and sleeping policemen, they are more focused and relentless. Humans are eclectic agents always engaged in a broad spectrum of activities. This makes them infinitely fascinating and adaptable, but also messes with their focus. Even the most dedicated professionals can only do their job for a few hours a day because they also need to eat, rest, talk, play, give and receive emotional support, and so on. Machines tend to be more specialized, which allows scaling up their efforts (the shoveling machine could not move so much earth, if it also had to pick flowers or play cello).

The keyword in the quote above is “inequality”. It captures the idea that the value of AI (as of any technology) comes from its alterity. This is an excellent example of Latour’s intellectual yoga: a posture that is both balanced and stretched outside the two commonplace positions we take when considering AI. On the one hand, the fear of a technology that could overpower us – as if there were not thousands of other technologies already overpowering us. On the other hand, the temptation to belittle AI by marketing it as a submissive assistant (as neo-Marxist scholars duly noted, casting AIs as inexpensive and non-unionized knowledge workers really speaks to the dullness of contemporary capitalism).

In other words, contemporary AIs are interesting not because they can succeed in Turing’s imitation game, but because they fail it productively. How do you know that a generative chatbot is not human? By its ability to answer questions on any knowledge domain, fielded by millions of users in dozens of different languages. No human could do that, and that’s precisely what makes AI intriguing.

4. AI as an actor-network

If, as discussed above, Latour was wary of the aim of the Turing test, he was even more skeptical of its conditions. Where Powers thinks that machines and humans may equally

cheat by mechanically emulating thought, Latour thinks that it is the test that is cheating by imaging a set up in which both the machine and the human are isolated and connected solely through a teleprinter:

The idea of a test matching a naked, isolated intelligent human against an isolated naked automated machine seems to me as unrealistic as imagining that we are here alone talking through email “naturally”, “directly”, without any mediation. Things and people are too much intertwined to be partitioned before the test begins, especially to capture this most heavily equipped of all faculties: intelligence. (Latour and Powers 1998, 180)

This critique of the Turing test resonates again with ANT. If the first half of this theory encourages us to regard everything that acts as an actor, the second posits that every action results from a complex network of interactions (Venturini and Munk 2021, 124-136). This is why ANT is also called “sociology of translation” (Callon 1986), because one of its core ideas is that no program of action can ever succeed when carried out by a single actor. Alliances and compromises are necessary to complete the smallest task (cf. Bijker et al. 1989 and Bijker and Law 1992). Discussing Hutchins’ (1995) book *Cognition in the Wild*, Latour argued that this is very much true of thinking as of any other type of action:

[...] cognition has nothing to do with minds nor with individuals but with the propagation of representations through various media, which are coordinated by a very lightly equipped human subject working in a group, inside a culture, with many artifacts and who might have internalized some parts of the process. (see Latour in Keller et al. 1996, 57)

and

Laboratories think, communities discover, disciplines progress, instruments see, not individual minds. (*ibid.*, 62)

Latour took this idea a step further. Since agency is defined by action, then the relational nature of actions implies the relational nature of actors. Latour explored this idea in many of his writings, but perhaps the clearest example comes from his reading of the controversy about gun control. Are guns neutral tools which can be used for good or ill depending on who wields them (as in the National Rifle Association’s slogan “guns don’t kill people, people kill people”) or are they inherently harmful instruments destined to realize their destructive function? Neither, answers Latour, guns have no essence separated from the persons wield them, but (and here is the real kicker) the same is true for their wielders:

Essence is existence and existence is action. If I define you by what you have (the gun), and by the series of associations that you enter into when you use what you have (when you fire the gun), then you are modified by the gun – more so or less so, depending on the weight of the other associations that you carry. This translation is wholly symmetrical. You are different with a gun in hand; the gun is different with you holding it. You are another subject because

you hold the gun; the gun is another object because it has entered into a relationship with you. (Latour 1994a, 33)

Hence the treachery of the Turing test. To pretend that there can be such a thing as an AI separated from the infrastructure that supports its existence, which includes many other machines as well as many other humans. And, symmetrically, that there can be such a thing as an equally isolated human tester. So much for the idea of the “singularity”, the hypothetical moment in which AI will break free from its human creators! If the history of technology has taught us anything, it is that nothing ever breaks free from anything. Quite the contrary, since prehistory and maybe even before, the evolution of humans and technologies is a chronicle of mutual entanglement and escalating interdependence (Latour 1993b; 1994b). The myth of the golem, the creature that emancipates and takes over its own creator has never been anything but a myth. Not because of the creature’s virtues or flaws, but because we, the creators, have always been a Frankenstein of our own. An actor-network, if you wish.

5. Iconoclasm, labor exploitation and the fabrication of AI

The misled obsession for AI’s autonomy recalls another fallacy that Latour questioned earlier in his career, that of iconoclasm – i.e., the belief that religious images cannot be authentic if they are fabricated. In his exhibition on the subject, Latour disputed the idea that acheiropoietic icons (not made by human hands but created miraculously) hold the highest value. Instead, he argues:

that the more human-work is shown, the better is their grasp of reality, of sanctity, of worship. That the more images, mediations, intermediaries, icons are multiplied and overtly fabricated, explicitly and publicly constructed, the more respect we [should] have for their capacities to welcome, to gather, to recollect truth and sanctity. (Latour 2002, 18, 20)

The exhibition extended this idea to science – claiming that the more and the better scientific facts are fabricated, the more they connect to their objects – and to politics, whose art lies precisely in stitching together an improbable community that would not exist otherwise. Here we see that Latour’s skepticism for autonomy and his interest for associations becomes normative. Not only he sees networks everywhere he looks, but he is also convinced that such a relational organization is *right and just, our duty and our salvation*.

If religious, scientific and political facts should be valued for their careful fabrication and their network of associations, why should AI be any different? In his conversation on AI, Latour connects the two ideas: “I don’t understand immaculate conception, more exactly I have great respect for the strange virginal dogma of the Church and cannot have any patience with its application to machines” (Latour and Powers 1998, 186). Only AI zealots can believe in the self-sufficiency of intelligent machines. We, humans, are intelligent machines and we are not self-sufficient at all. Examine any modern city and you will witness an incredible infrastructural effort necessary to bring in the food, water, heat we need and bring out all the waste that we

produce. And that is just what's needed to keep us alive (Latour and Hermant 1998). If, on the top of that, you also want humans to be able to think, then be ready to add transportation, electrical and communication cables, payrolls, education systems, museums, libraries, theaters... and, of course, all the political, social and administrative institutions necessary to manage them.

The same is obviously true for AIs. According to some observers, the recent step forward in AIs capacities comes from the alignment of a plurality of elements of different nature (Huang 2023): social (the explosion of user-generated content used to train the algorithms); cultural (the shift from symbolic to connectionist AI); technological (the repurposing of graphics processing units); mathematical (the introduction of “transformers” with capacity for parallelization and self-attention); financial (low interests rates allowing to invest in experimental project) etc. Thanks to this heterogeneous network of allies, contemporary AI can do things that were unimaginable for their less plugged-in predecessors. And that's not all, by flipping the myth of the self-governing AI on its head, generative AIs are amazing not because they are autonomous, but because they are wired to a vast socio-technical infrastructure.

This dispels the fuss around deep and unsupervised learning. While these expressions signal an important discontinuity from earlier approaches, they do not mean that AIs are capable of learning by themselves. They just mean that human intervention has been displaced. In the older “symbolic approach”, experts were consulted to formulate explicit rules to be fed into the machine. In the current “connectionist approach”, the machine is given a very long series of examples and ask to detect patterns in them (e.g., that Latourian texts tend to revolve around words such as networks and associations, while Bourdieusian texts around as power and capital). This learning is said to be “deep” (or black-boxed) because the algorithm generated by the learning does not have to be explicit or understandable by a human. It is said to be “unsupervised” if the examples have not been tagged by human coders (e.g., no one told the computer which texts were by which author, it discovered this by noticing their signatures).

Clearly neither “deep” nor “unsupervised” mean autonomous. Quite the contrary, both require the help of a small army of mathematicians, data scientists, engineers, user experience designers, system managers, company managers, venture capitalists, etc. Commenting Deep Blue victory against Kasparov, Latour (1997) observed:

They say: homo sapiens against the machine. Quickly said. Rather, it's homo sapiens in one form (world chess association, Kasparov, hundreds of years of gaming tradition) versus homo sapiens in another form (chess games from throughout history in memory, the millions of hours of work accumulated by hundreds of IBM programmers, the hundreds of years of experience in PLCs).

Also, machines need vast volumes of examples to learn from, which need to be generated by humans. Some algorithms are trained by “adversarial learning”, when two or more machines are asked to play against each other to improve their performances, but this only works for tasks (such as security attacks or chess games) whose success is easy to define. While there are clear rules to determine whether a chess game lost, it is difficult to say what an interesting text reads like – so you cannot just let two chatbots discuss to improve their conversation. Instead, you need a huge army of humans producing contents that the machine can learn to imitate.

Machine learning is only as good as the data it learns from. This is why standard datasets such as *ImageNet*, *BookCorpus*, *WikiText*, and others are keystones to AI (Denton et al. 2021). Obviously, none of these datasets is autonomous, unsupervised, or self-assembled. Quite the contrary, all of them have required a massive human investment for the production, harvesting, cleaning, and labeling of their data (Pasquinelli 2019). These activities have often been crowdsourced using micro-labor platforms exploiting workers in the Global South (Tubaro et al. 2020) or ripping the Web without the slightest recognition or remuneration. None of the fan-artists filling the galleries of *Deviant Art* or *AO3*, the contributors editing the entries of *Wikipedia*, the developers providing advice on *Stack Overflow*, and in general none of the users posting content online did so with the intention of satiating AI's appetite for examples. Some artists are beginning to file lawsuits against the unauthorized and uncompensated use of their work, but their chances are uncertain. No company in the world has the money to pay for the gigantic amount of content fed into AI training datasets and courts and governments may be hesitant to cripple this nascent industry. Not to mention that the cloud-based nature of this technology makes it easy for companies to move to less regulated countries. Talk about barking at the wrong tree! We worry that AI might emancipate from its creator and forget that AI companies are so dependent on our input that they are ready to steal it off from us.

As Latour argued in many of its writings, it is the individual machines that we should worry about, but the sociotechnical systems in which they participate. This understanding radically changes our means of resistance. If we fear the menace of a rogue machine, then the best we can do is try to unplug through math or engineering (as in many sci-fi movies). If we oppose an industrial conglomerate, then our weapons are economic incentives, legal provisions, and societal mobilization (admittedly a lousy sci-fi movie!). *Suivez le réseau!* ("Follow the network!") would say Latour, and his idea has not lost a bit of relevance.

6. AI's mode of existence

So far, I have reflected on AI in the light of the general Latourian sociology of technology, yet one can also use his philosophy to pinpoint how generative AI is different from other technologies, its distinctive "mode of existence" as Latour would say. Here, Latour can be useful not only to social scientists but also to data and computer scientists.

In *An Inquiry into Modes of Existence* (2013) Latour suggests that, in order to understand the nature of beings, we need to examine how they act or, more precisely, the "felicity conditions" allowing to judge whether their action is successful. This idea comes of course from the philosophy of language (Austin 1962; Searle 1969), but Latour gave it a more existentialist meaning. Saying that the scientific method revolves around building reversible chains of reference means that scientific facts exist only as long as a connection between natural phenomena and their description is maintained. Lose this connection and your fact goes poof.

The felicity condition of technology in general is the replacement of human actions by non-human ones through what Latour calls: "the technological detour... [that] makes it possible, not to do something, but to *have something done*" (Latour 2013, 229, *emphasis in the original*). Crucially, however, since this is a detour and not a simple delegation, the thing that

used to be done is now done differently. An excavator does not just move more earth than a human, it also moves it differently, just as Hollywood moves emotions differently from a human storyteller. This also applies to AI, which is interesting precisely because its way of drawing and writing is different from that of humans. But different *how*?

Here, I believe, we should differentiate between the experimental development of generative AIs and their current and near-future industrial applications. Technologies in the laboratory are not the same thing as technologies on the market (Latour 1987). In the laboratory, generative AI is judged through what is called “reinforcement learning”, that is by asking human evaluators to judge AI’s answers. The need for such external validation is far from obvious. If we believe in the idea of a self-teaching AI, why should we add yet another army of humans, this of judges, on top of those that already produced, harvested and cleaned the examples the machine learns from? Why not let the machine establish its own felicity conditions?

Interestingly, this was a pragmatic call. Reinforcement learning was introduced after early AIs trained with content pulled from the Web ended up embracing racist and chauvinist stereotypes and, in some cases, extreme hate speech. This shouldn’t have come as a surprise: the Web can be a toxic place and many of the training datasets were laden with biases (Crawford and Paglen 2021; Bandy and Vincent 2021). Thus, the dream of AI’s autonomy was once more betrayed and more human connections were added to make the learning a bit less automatic. Incidentally, this has been noted and objected by right-wing groups that believe that their hard-fought online influence should not have been reinforced away. And these groups are now engaged in developing new AIs more strongly associated with their values. Once again, the problem is not AI’s autonomy, but its associations.

If human reinforcement defines the felicity condition of AI *in the laboratory*, nothing assures that it will keep defining it *in the market*, where, I am afraid, felicity conditions are quite different. Indeed, if the Web is often a toxic place it is not only because it is full of toxic people, but also because the infrastructures that handle the circulation of online contents are, in many ways, partial to toxic content. To understand why, one needs to remember that the business model of most online platforms is advertising. Contemporary digital technologies are rooted in the attention economy of digital marketing and, in this economy, content that is outraging, hyper-partisan, sensationalist or otherwise attention-grabbing is highly favored (Venturini 2019). As candidly admitted by one of the creators of the neural-network recommendation algorithm used by YouTube:

In addition to the first-order effect of simply recommending new videos that users want to watch, there is a critical secondary phenomenon of boot-strapping and propagating viral content. (Covington et al. 2016, 193).

Sure, generative artificial intelligence is not only used for tasks related to online attention economy, but consider this: it is trained on content generated by this economy; it is financed and dominated by the same companies that rules this economy; and generally, more and more social activities are migrating on the technologies (phones and computers), infrastructures (the Web and the Internet) and metrics (likes and shares) that supports this economy. What do you think will likely happen?

What is already happening is that generative AIs are being used to produce contents whose felicity is ultimately evaluated not by some human *arbitrer elegantiae*, but by a sociotechnical system composed, human influencers, social media platforms, online advertisers *and* deep learning recommendation algorithms. In line with McLuhan's prophecy: the "message" of all new media and technology is its way to "amplify or accelerate existing processes" (1964, 8). We will soon be in a situation in which generative AIs create content that humans pick and share, hoping to please recommendation AIs and thus be widely circulated online, thereby gaining a higher probability of influencing the further AI training. Mind you, it is not the circularity of the process that worries me (this is not another phony singularity argument); it is the fact that this cycle risks being steered by the very dull felicity conditions of maximizing our addiction to digital advertising. God forbid!

7. Conclusions

This scenario explored generative AI through the lens of Bruno Latour's philosophy of technology. It started by listing a few reasons why Latour did not seem particularly interested in AI. I noted that the project of simulating human intelligence feels utterly out of sync with its latest interest in the ecological crisis and its political consequences. I observed that the opacity of machine learning is at odds with Latour's definition of science as a reversible chain of reference. And I remarked that, as an ethnographer and someone obsessed with details and outliers, he would have disliked the black-boxed aggregation of deep learning. If anything, Latour might have appreciated AI as a writing tool, an instrument capable of quite subtle language games, a skill that he saw as crucial to research in social sciences.

Apart from what Latour might have liked or disliked, a Latourian perspective helps bring into focus the threats and promises of AI. Unlike what the Turing test or the idea of the singularity suggests, AI is not separate or separable from the vast network of technological, financial, and organizational supporters that maintain its existence. Far from being autonomous, AI relies on a vast army of human allies, which includes the scientists and engineers who tweak its models, but also the people who produce and prepare its training dataset and who correct the relevance of its responses. Rather than fearing that AI might break free from us, we should worry about how AI is embedded in the infrastructures of the digital attention economy and aligned to the goals of digital advertising.

Yet, not all hope is lost. One of the most important things that I have learned from Latour is that the dynamics of science and technology are always bound to defy our expectations. Resulting from a multiplicity of unpredictable interferences, the technological detour always takes us to places we did not anticipate. This cast an interesting light on the much-discussed issue of "AI's hallucinations": the tendency of generative machines to *extrapolate* from their training on and offer outputs that, though plausible, are factually false and sometimes wildly delusional. In the early stages of generative AI, this was considered a fascinating feature – so much so that one of the first image generators was proudly named *DeepDream*. Yet as the technology matured, hallucinations have come to be considered a source of misinformation and a threat to public debate.

AI fantastic extrapolations are indeed a serious problem, but only if we accept that AI can serve as a reliable substitute for newspapers and text-books or, even worse, for journalists and school teachers. As Latour taught us, such an idea is both false and dangerous. No technological being can ever serve as a neutral substitute for another being (be it technical or human), and technological replacement always comes with a profound transformation of the situation, that it is better to understand than to deny. Rather than fighting AI hallucinations (in the misplaced hope that AI can be turned into a cheap substitute for intellectual labor), we should cherish its capacity for creativity. This is exactly what Latour did, drawing a parallel between machine intelligence and critical thinking:

What would critique [and AI] do if it could be associated with more, not with less, with multiplication, not subtraction. Critical theory died away long ago; can we become critical again, in the sense here offered by Turing? That is, generating more ideas than we have received, inheriting from a prestigious critical tradition but not letting it die away, or “dropping into quiescence” like a piano no longer struck. This would require that all entities, including computers, cease to be objects defined simply by their inputs and outputs and become again things, mediating, assembling, gathering many more folds... Then we would have gone for good beyond iconoclasm. (Latour 2003b, 248)

References

- Austin, John L. (1962) *How to Do Things With Words*, Oxford, Clarendon Press.
- Bandy, Jack and Vincent, Nicholas (2021) *Addressing “Documentation Debt” in Machine Learning Research: A Retrospective Datasheet for BookCorpus*, in “ArXiv: 2105.05241”, pp. 1-13.
- Bareis, Jascha and Katzenbach, Christian (2022) *Talking AI into Being: The Narratives and Imaginaries of National AI Strategies and Their Performative Politics*, in “Science, Technology, & Human Values”, 47(5), pp. 855-881.
- Bijker, Weibe E., Hughes, Thomas P. and Pinch, Trevor (eds.) (1989) *The Social Construction of Technological Systems*, Cambridge (MA), The MIT Press.
- Bijker, Wiebe E. and Law, John (1992) *Shaping Technology/Building Society: Studies in Sociotechnical Change*, Cambridge (MA), The MIT Press.
- Callon, Michel (1986) *Some Elements of a Sociology of Translation: Domestication of the Scallops and the Fishermen of St Brieuc Bay*, in John Law (ed.), *Power, Action and Belief: A New Sociology of Knowledge?*, Boston, Routledge and Kegan Paul, pp. 196-233.
- Cardon, Dominique, Cointet, Jean-Philippe and Mazières, Antoine (2018) *Neurons Spike Back*, in “Réseaux”, 211(5), pp. 173-220.
- Coccia, Emanuele (2021) *An Interview with Bruno Latour on The Art of Writing*, in “Asymptote”. Available at: <https://www.asymptotejournal.com/interview/an-interview-with-bruno-latour/> (retrieved on October 27, 2023).
- Covington, Paul, Adams, Jay and Sargin, Emre (2016, September) *Deep Neural Networks for YouTube Recommendations*, in “RecSys ‘16: Proceedings of the 10th ACM Conference on Recommender Systems”, Boston (MA), pp. 191-198.

- Crawford, Kate and Paglen, Trevor (2021) *Excavating AI: The Politics of Images in Machine Learning Training Sets*, in "AI and Society", 36(4), pp. 1105-1116.
- Denton, Emily, Hanna, Alex, Amironesei, Razvan, Smart, Andrew and Nicole, Hilary (2021) *On the Genealogy of Machine Learning Datasets: A Critical History of ImageNet*, in "Big Data & Society", 8(2).
- Ferrari, Fabian and McKelvey, Fenwick (2023) *Hyperproduction: A social theory of deep generative models*, in "Distinktion: Journal of Social Theory", 24(2), pp. 338-360.
- Huang, Haomiao (2023, January 30) *The generative AI revolution has begun – how did we get here?*. Ars Technica. Available at: <https://arstechnica.com/gadgets/2023/01/the-generative-ai-revolution-has-begun-how-did-we-get-here/> (retrieved on October 27, 2023).
- Hutchins, Edwin (1995) *Cognition in the Wild*, Cambridge (MA), The MIT Press.
- Katzenbach, Christian (2021) "AI will fix this" – *The Technical, Discursive, and Political Turn to AI in Governing Communication*, in "Big Data & Society", 8(2).
- Keller, Janet Dixon, Bazerman, Charles and Latour, Bruno (1996) *Cognition in the Wild (Book)*, in "Mind, Culture, and Activity", 3(1), pp. 46-63.
- Latour, Bruno (1987) *Science in Action: How to Follow Scientists and Engineers Through Society*, Cambridge (MA), Harvard University Press.
- Latour, Bruno (1993a) *La clef de Berlin et autres leçons d'un amateur de sciences*, Paris, La Découverte.
- Latour, Bruno (1993b) *We Have Never Been Modern*, Cambridge (MA), Harvard University Press.
- Latour, Bruno (1994a) *On Technical Mediation: Philosophy, Sociology, Genealogy*, in "Common Knowledge", 2(3), pp. 29-64.
- Latour, Bruno (1994b) *Pragmatogonies: A Mythical Account of How Humans and Non-Humans Swap Properties*, in "American Behavioral Science", 37(6), pp. 791-808.
- Latour, Bruno (1995) *The "Pédofil" of Boa Vista: A Photo-Philosophical Montage*, in "Common Knowledge", 4(1), pp. 144-187.
- Latour, Bruno (1997) *Rien d'étonnant à ce que la machine fasse mieux*. Liberation. Available at: https://www.liberation.fr/livres/1997/05/13/rien-d-etonnant-a-ce-que-la-machine-fasse-mieux-entretien-avec-le-professeur-bruno-latour_205823/ (retrieved on October 27, 2023).
- Latour, Bruno (2003a) *The Promises of Constructivism*, in Don Ihde and Evan Selinger (eds.), *Chasing Techno-science: Matrix for Materiality*, Bloomington, Indiana University Press, pp. 27-46.
- Latour, Bruno (2003b) *Why Has Critique Run Out of Steam? From Matters of Fact to Matters of Concern*, in "Critical Inquiry – Special issue on the Future of Critique", 30(2), pp. 225-248.
- Latour, Bruno (2005) *Reassembling the Social: An Introduction to Actor-Network-Theory*, Oxford, Oxford University Press.
- Latour, Bruno (2007, April 6) *Beware, your imagination leaves digital traces*, in "Times Higher Literary Supplement". Available at: <http://www.bruno-latour.fr/sites/default/files/P-129-THES-GB.pdf> (retrieved on October 27, 2023).
- Latour, Bruno (2013) *An Inquiry Into Modes of Existence: An Anthropology of the Moderns*, Cambridge (MA), Harvard University Press.
- Latour, Bruno (2014) *L'influence est un risque*, in Claire Tollis, Laurence Créton-Cazanave and Benoit Aublet (eds.), *L'effet Latour: Ses modes d'existence dans les travaux doctoraux*, Paris, Editions Glyphe, pp. 1-5.
- Latour, Bruno and Hermant, Emilie (1998) *Paris Ville Invisible*. Paris: La Découverte-Les Empêcheurs de penser en rond.

- Latour, Bruno and Powers, Richard (1998) *Two writers face one Turing test: A dialogue in honor of HAL*, in "Common Knowledge", 7, pp. 177-191.
- Latour, Bruno (2002) *What Is Iconoclasm? Or Is There a World Beyond the Image Wars?*, in Bruno Latour and Peter Weibel (eds.), *Iconoclasm: Beyond the Image Wars in Science, Religion and Art*, Cambridge (MA), The MIT Press, pp. 14-37.
- Latour, Bruno, Jensen, Pablo, Venturini, Tommaso, Grauwin, Sébastien and Boullier, Dominique (2012) *"The whole is always smaller than its parts" – A digital test of Gabriel Tarde's monads*, in "The British Journal of Sociology", 63(4), pp. 590-615.
- McLuhan, Marshall (1964) *Understanding Media: The Extensions of Man*, New York, McGraw-Hill.
- Pasquinelli, Matteo (2019) *How a Machine Learns and Fails: A Grammar of Error for Artificial Intelligence*, in "spheres: Journal for Digital Cultures", 5, pp. 1-17.
- Searle, John Rogers (1969) *Speech Acts: An Essay in the Philosophy of Language*, Cambridge (UK), Cambridge University Press.
- Suchman, Lucy (2008) *Feminist STS and the Sciences of the Artificial*, in "The Handbook of Science and Technology Studies", 3(1), pp. 139-163.
- Tubaro, Paola, Casilli, Antonio A. and Coville, Marion (2020) *The Trainer, the Verifier, the Imitator: Three Ways in Which Human Platform Workers Support Artificial Intelligence*, in "Big Data & Society", 7(1).
- Veale, Michael, Matus, Kira and Gorwa, Robert (2023) *AI and Global Governance: Modalities, Rationales, Tensions*, in "Annual Review of Law and Social Science", 19(1), pp. 255-275.
- Venturini, Tommaso (2019) *From Fake to Junk News, the Data Politics of Online Virality*, in Didier Bigo, Engin Isin and Evelyn Ruppert (eds.), *Data Politics: Worlds, Subjects, Rights*, London, Routledge, pp. 123-144.
- Venturini, Tommaso (2024, forthcoming) *Quali-Quantitative Methods*, in Alan Irwin and Ulrike Felt (eds.), *Encyclopedia of Science and Technology Studies*.
- Venturini, Tommaso and Munk, Anders (2021) *Controversy Mapping: A Field Guide*, Cambridge (MA): Polity Press.
- Venturini, Tommaso, Jacomy, Mathieu, Meunier, Axel and Latour, Bruno (2017) *An Unexpected Journey: A Few Lessons from Sciences Po Médialab's Experience*, in "Big Data & Society", 4(2), 205395171772094.
- Venturini, Tommaso, Jensen, Pablo and Latour, Bruno (2015) *Fill in the Gap: A New Alliance for Social and Natural Sciences*, in "JASSS – Journal of Artificial Societies and Social Simulation", 18(2).