

The magical relationship to AI, or how to live with technical alienation

Tyler Reigeluth

Tyler Reigeluth holds a PhD in philosophy from the Université Libre de Bruxelles, within the “Algorithmic Governmentality” research project. He is assistant professor of philosophy at the Catholic University of Lille, at the ETHICS research center. His research focuses on the normative interactions between human and machine learning as well as questions surrounding algorithmic and technical transparency. He co-authored *Éthique de l’information et de la communication* with Thomas Berns (Presses Universitaires de Bruxelles, 2021) and *Villes intelligentes. Critique d’une transparence sans fin* (Météores 2023).

T.C. In your presentation, you mentioned the fetish of automation in Western culture. Could you give us a quick overview of its roots and evolution? What makes the relation between the West and automation special?

T.R. One of the reasons I speak about the fetish of automation in Western culture is to underline the fact that there's also a relationship of desire to our technical objects and to technical performance. We expect certain things of technological devices that aren't necessarily related to how they work concretely but to the expectation we have of how they should behave. And there's this a recurring idea—at least since the beginning of the Industrial Revolution—that machines can, should and will replace human labor for reasons of efficiency: they can do things faster and with less error; they don't grow tired or go on strike; they aren't sick or they don't try to protest for their labor conditions. So there's this idea that machines are going to be able to replace and supplant human labor and even human intelligence, which is I think what we're seeing today with a lot of discourses around Artificial Intelligence. There is this largely unquestioned desire and idea that they can, and someday will, behave absolutely automatically that is to say without any human interference.

I think the reason why there's a sort of special relationship in the West, is because it's where the Industrial Revolution was born and where there is a very clear distinction between manual and intellectual labor, between those who have ideas and those who execute, between masters and slaves, managers and employees. There's been a lot of critical literature, namely in postcolonial theories of technology that shows the sort of continuity between the master/slave metaphor and the language used to talk about technology and more specifically about computer science. We sort of projected onto technology, in our Western culture, this idea that the only purpose of technology is to obey, to be sort of the ideal worker.

The birth of intellectual property is related to this conception that labor and the idea that labor executes can be clearly distinguished in a unidirectional relationship. this distinction usually tells us that there is no more intelligence in labor, it is just this sort of blind repetitive execution of a more abstract idea and it's actually idea or the production process that you patent. it's important to underline this because it shows that it isn't technology that turns workers into machines. Workers

were already turned into machines as much as they could be throughout the Industrial Revolution. And it's because they were that they could be replaced with machines. So we first had to de-class labor into a purely mechanical process. If chatbots today are used in client-service relationships it's because the client-service relationship was already nearly automated through incredibly scripted and controlled human interactions.

But things began shifting already before the Industrial Revolution. Around the 16th-17th century the idea emerged that the human body could be studied as a machine and both were ultimately regulated by the same sort of physical laws. The more we know those laws the more value and efficiency can be extracted from that "human machine". And throughout the Enlightenment, there's this paradoxical moment where the European intelligentsia is both very interested in the prospects of political emancipation but also fascinated with reproducing the physics of the human body, with and within automata. But it's quite striking to see that those automata are always either women, animals or Orientalized figures like the mechanical Turk; it's never a white man, who is the one endowed with the power of free thought. Thought is already presented as a power *over* something or someone, less free, less rational, more mechanical.

A lot of my thinking about the relationship between humans and technology, and AI more specifically revolves around the question of what it would be like if this relationship did not rely upon such a fetish of automation, this desire for technology to do what we ask it to while also replacing us in what we do. A counter-example would be the hacker or low-tech cultures, people who go to landfills or develop networks to salvage components that can be used to fix objects or for aesthetic or educational purposes. Those are relationships with technology that see the value of objects beyond simply their instrumental reality and what Gilbert Simondon called their *abstract function*. Consider for example a toaster: it is made to toast bread on that's it's function, that's how it's sold. And as soon as it doesn't do that then we get rid of it because it doesn't fulfill that abstract function. But on a very concrete material level they're actually many of aspects and components of it that still work very well and that could be reused in another object or maybe the object could be used to do something else, etc.

T.C. Following the Enlightenment, one could have thought that, as knowledge spread,

as science and technics progressed, magical belief would be challenged and ultimately, bound to disappear. But each technical invention actually comes with its own myths and superstitions. How would you say technic and magic are related?

The relationship with Modern techniques—and researchers have a synoptic view of what it is the due in the larger production and innovation process.

T.R. mean techniques that started evolving from the 16th century onward—within the overarching Cartesian framework of the separation between a mind and a body is that the machine’s mechanism was nothing else but the materialization of what it’s inventor had thought up. Thus the idea that technology was somehow going to free us from magical thinking. We were finally going to master the secrets of nature, of the human body and of all these things for which we lacked an analytical representation. Explaining something mechanistically means being able to follow it every step of the way. But this understanding of technology works for a very limited number of objects and for most of the users of those objects there is still a very magical dimension to how even the most quotidian objects behave. And that’s again been amplified since the Industrial Revolution, since we’ve started developing a mass production and consumption of these objects. As consumers and passive users of the object we are supposed to interact with them without understanding their concrete functioning. In no way are the users supposed to take part in the (re)invention or design of the objects they use.

If you look at most of the domestic objects in use—whether it’s toasters or computers—they are encased. Their inner functioning is hidden from us. And so it’s hard for contemporary users to have any relationship with their objects other than magical. We are perhaps far more alienated from our technical objects today than we were a thousand years ago, or five hundred, or even a hundred years ago when the mechanism was always, at least in part, visible and part of the fascination was to be able to see the mechanism at work. Think of windmills or locomotives for example. Which doesn’t mean that everyone would understand perfectly how they worked of course, but it could give a sense of intuition, a feeling and also some aesthetic pleasure by watching the mechanism work. Today we press on buttons and then something happens. And it’s rather fascinating to see how a very limited number of interfaces have replaced the incredible variety of workings behind those interfaces.

This magical thinking also affects engineering communities and not just the users. The social division of labor also concerns form of labor that are traditionally considered more abstract or intellectual, including large scale techno-scientific projects and companies where very few employees

Another aspect, more specifically related to AI research, is that a lot of the algorithms and the technical systems function at a scale and a complexity that baffles individual competences. The behaviors of many technical systems are increasingly non-deterministic. They are partly emergent—meaning that their behaviors emerge through the learning process itself. They weren’t written out or explicitly programmed ahead of time. They are highly probabilistic so that makes it very hard for engineers who are trained in a culture of control to adapt their epistemological frameworks and communicate to the larger public when things don’t go according to plan. This happens for example when engineers at Google or Microsoft can’t exactly account for what went wrong in a machine learning application that produces racist or sexist outputs.. They can’t say “OK, it’s this step in the algorithm where it all went wrong”. They can only make hypotheses and then try to tweak or tune the algorithm and see if it changes its behavior.

I find this fascinating that an engineer is confronted to that sort of behavior that I think a lot of end users already feel well before the advent of AI. Most of us didn’t have to wait for AI to feel like we’re surrounded by magic objects.

T.C. Can you explain what Machine Learning is and how it breaks with the classical paradigm of algorithmic programming?

T.R. Given its recent success today Machine Learning is more or less conflated with AI, but historically it is a subfield of AI. It tries to develop algorithmic systems that can extract a model from the data rather than simply apply in a model to the data. In supervised learning, the machine is given labeled data and told “this is a cat, this is a cat, that’s a dog, that’s a dog”. It’s done thousands of times—that’s human labor in supervised learning systems by the way—there is nothing automatic about that. And then the algorithm develops its own abstract model of what a cat and dog are.

What's important to underline is that based on that abstract model it's going to make predictions after, on new data that hasn't seen, on new pictures for example of cats and dogs. It is probabilistic, not deterministic. There's always a level of uncertainty in its behavior. Which is why those kinds of algorithms are very good in dynamic interactions, like the kinds that are used in recommender systems on Spotify, YouTube or Netflix. "I'm 80% sure—Netflix says—that you're going to like this music." The recommendation process is very open to the fact that you won't like it and that's okay; it will just suggest something else and fine tune the probability of the recommendations based on the interactions with the user. The recommendation process is resembles a sort of iterative relationship or dance between the algorithms and the users.

I think it's important to emphasize that these machines are doing some parts of the job automatically but they aren't learning automatically. Learning is a social activity; it's not something you do by yourself without any contact with anybody. If we're very serious about machines learning, we may need to think about educating them. Many of the problems around machine biases, that are racist, xenophobic or misogynist surprise us because we have this idea that we can train them on data in a lab and then have them sort of develop good models "out there". It's rather naive and reductive understanding of how learning actually works.

We should consider training as a more important form of labor than just something that should be outsourced and the social contexts from which data is extracted as more than just a resource. Alan Turing, in the 1940's, was already saying that it would be rather unfair and misguided to compare a machine who has to learn things very quickly in a lab with a human who spent eighteen or more years of its life learning and very rich social interactions. Today, we're surprised that algorithms are potentially racist or xenophobic. To me that's not the problem in and of itself. The problem is that we're surprised the machines develop those biases. We want them to learn by themselves while also giving us objective results. So there's a real tension there that's quite revelatory of where we are at in a relationship with technology.

T.C. Arts and crafts turned into technology and took over a large spectrum of human operations. This transformation has raised concerns—among them the idea of "technical alienation", declined in many ways in

the Western political thought. What is Simondon's proposal for "technical culture" and how is it a different take at addressing technical alienation?

The point I'm trying to make, based on **T.R.** Simondon's idea of technical culture is that it's not just a question of individual know-how, it's not just a question of opening the black box, but it's also question of having the right institutions and forms of education that foster uses that aren't the ones that are expected by their engineering and design. The user should actually be able and allowed to invent new ways of using a technical object. This involves theoretical knowledge but also a situated, manual, sensitive and embodied experience of those objects, because you can have a very good understanding of a machine without necessarily having a theoretical understanding of it. You can understand how your lawnmower works because you spent a lot of time with it and you know that when it vibrates at a certain place then you have to change the lever or do this or that without actually being a mechanic. Now, I think we need to move beyond Simondon a bit, however, because the complexity and scale of the technical systems we use today, especially computational and AI systems are not analogous to individual know-how and mastery. The ways in which we can overcome our technical alienation are necessarily going to rely on forms of collective assemblages (what Simondon would call collective individuation). How do we organize people with different skills and forms of knowledge into a collective that can together surmount that technical alienation? I think a good example of that on a local scale are things like repair cafés where a network of people with different forms of knowledge, different tools, traditions and concerns can think, repair and take care of things collectively. The solution our society tends to offer us is simply to call an expert who costs a lot of money and who really only puts an expensive band-aid on our alienation because we rarely learn anything from this interaction that would enable us to better understand or master our technical object.

It's worth noting that many of the algorithmic systems we use already involve that kind of collective intelligence. The difference is that they try to hide it from the user. Systems like Spotify, YouTube or Netflix rely on rich social interactions. There's a whole aspect of their interaction that is collective and not individual at all, but that's hidden behind interfacing and design choices that make the interaction all about "personalization"

and targeting.

Looking beyond the interface, not taking things at “(inter)face value” as Sherry Turkle put it, is often seen as problematic for the platforms developing the algorithms who see such users as trying to “game the algorithm”. Such users are able to understand or intuit the way the algorithm works without actually accessing the source code or opening up the black box. They can develop a certain understanding by playing with the algorithm, in the very noble sense of the term, by trying to predict its inner workings through interaction, through repeated use, trial and error.

It’s exactly what happened with Microsoft chat bot Tay back in 2016. It was taken offline after several hours because it started developing anti-semitic and racist slurs. And what’s fascinating is that Microsoft said “ill-intentioned users tried gaming and exploiting a vulnerability in the algorithm”, implicitly saying there are ways the users should and shouldn’t use what we give them. To me, that’s exactly what Simondon is talking about in terms of technical culture and the attempts by these platforms to limit the user’s natural right to game the algorithm, shows us the degree to which it’s not about the user at all but it’s about profit extraction and other forms political control, economic competitiveness and intellectual property.

T.C. During your presentation, you introduced Gilbert Simondon’s notion of “margin of indeterminacy”. Could you explain what it means? How is this notion relevant for our use of and relation to Machine Learning?

T.R. His idea of margin of indeterminacy is counter intuitive for us because he basically says that the more complex the machine is, the more opportunities it affords for interactions and the more unpredictable it becomes. So a hammer for example has a very low margin of indeterminacy because it’s structure is given once and for all. A clock has a bit more than a hammer but basically has a very low level of indeterminacy because it doesn’t receive any outside information. Machines that have a margin of indeterminacy are machines that communicate with their environment, like windmills for example. But for the windmill that’s relatively controlled and limited: It has a certain kind of input it can receive, at a certain place and it’s behavior is highly determined

Machines that have integrate a self-regulatory function present a much higher margin of indeterminacy. For example, a steam engine has a reg-

ulatory system within it to maintain the pressure at a certain level, because it has information coming in, which is the accumulation of pressure in its compression chamber, and has to keep that pressure at a certain level, using a governor.

Simondon is interested in complex machines that regulate their behaviors through internal but also external mechanisms. Internal mechanisms are built into the machine and allow for example to regulate it’s pressure just like the human body has an internal regulator that keeps its temperature at the same level. External mechanisms are the ones where the human intervenes to regulate the machine’s functioning, to make sure it’s operating within its viable norms; just like we might open the window or adjust the thermostat if we are hot.

So what he’s trying to point out is that the more complex the machine, the more opportunities there are for us to interact with it. In a sense, he’s taking the reverse position on the history of automation : it’s not that machines are becoming more and more automatic; aspects of their behavior are more automatic for sure but that doesn’t mean that they do everything absolutely automatically. Establishing where the margin of indeterminacy is, is something that is essential to developing the technical culture I was talking about earlier...

A recent example is an artist-activist in Germany who wanted to ridicule the smart city project in his local town and the presence of Google namely how it’s services are used to traffic management. He loaded a wheelbarrow with 99 geolocalized smartphones and it created a traffic jam on Google Maps. That’s a good example of gaming the algorithm by understanding that it actually has a margin of indeterminacy. There’s space there for you to disrupt the logic of its function and what’s interesting is that it doesn’t mean that you’re going to make it dysfunction, you aren’t breaking the machine by exploring its margin of indeterminacy. There are multiple implicit purposes and that often what is sold as the sole purpose is actually the company’s purpose.

Maybe we all have a fundamental right to use technology as we want. Perhaps that’s a more provocative or political thing to say, but I think that’s something worth thinking about today as we’re increasingly surrounded by technological mediations of all kinds, be they “intelligent” or dumb.

T.C. One of the reasons, you said, programmed obsolescence works today is because we fail at recognizing value in past old technological objects. Could you elaborate

on that?

T.R. Simondon already anticipated this idea of programmed obsolescence of objects, and that's why he talked about a certain level of magic that was involved in the fetish of automation. If you look at the advertising in the 50's, 60's and 70's, it became about selling a magical experience of technology, as something that was going to free you up to do other things. We don't actually buy technical objects anymore, we are buy objects of consumption. They become generic whether it's a toaster or a car, it doesn't really matter: what we're buying is this feeling of novelty, of magic when it operates by itself, when all you have to do is press a button.

An important aspect of Simondon's notion of technical culture is that any given object has its specific history, its own evolution and that understanding that evolution can help us reappropriate the object but also ourselves insofar as human history goes hand in hand with technical history. For example, understanding the Middle Ages requires us to understand how their technical objects work: their mills, their building cranes, their cooking utensils... and if we lose that we lose also a continuity in cultural and social history that connects us with the past. And we see that quite cruelly today in the innovation economy. Inventions just seem to come out of nowhere; they have no continuity in history and so they also have no value because they're disposable. So when Simondon says there's a value in old objects he's is not trying to replace one fetish with another but he's pointing out that it will make our technological future richer, more balanced. There is a kind of respect for technical objects and operations that needs to be fostered.

Programmed obsolescence works today because the users have little culture of their objects. When the printer breaks, it breaks. At no point can we look at it and say "oh that looks kind of like a piece that's on my washing machine, maybe I can actually use that piece in here and try that there". And they're probably lots of instances where you can do that actually, where they are relatively generic interchangeable pieces or mechanisms that are suitable and it would work well enough. Actually, not that many when it comes to mass produced objects because their design is made to avoid any such participation, creativity or inventiveness by the user.

T.C. During your lecture, you mentioned machinal behavior as being first a human characteristic, through the con-

cept of *mental penombra* developed by anthropologist Leroi-Gourhan. You also said that the cloud could be considered as a "new technological unconscious" in the sense that "it's doing things for us in the background". Could you please explain those ideas?

T.R. More recently I have been interested in trying to understand how a large part of those distributed, multi-scalar systems we use today can be considered in what the paleo-anthropologist André Leroi-Gourhan called a *pénombre psychique* [mental penumbra]. In his understanding of our interactions with technology, he distinguishes three different levels of behavior: lucid, machinal and automatic behaviors. All three of these are human before being mechanical, that is performed by an objective machine.

The automatic level is basically our motor reflexes: things we can barely control like breathing or blinking. Machinal behaviors, which make up the majority of our behaviors, are learned socially by imitation, trial and error, and communication. They are the fabric of our cultural habits. Once we've learned them correctly, we perform these behaviors without thinking about them, but we can always foreground them when they things don't go according to plan. I don't think about opening the door but if I can't open the door all of a sudden then opening it becomes a lucid or conscious behavior. The conscious or lucid behaviors are ones that we've been engaged in for the past hour here, deliberately and consciously thinking about what we're doing. Those are quite rare though in our daily life.

This distinction helps us understand how we've come to share certain behaviors with complex machines and technical systems that are adapting their behaviors in relation to ours; they're interacting with us, they're not simply executing orders. And that part of our machinal behaviors are actually now also embedded in the technical system. A good example of this is an escalator coming out of a metro station. Walking up stairs is something we learned, it is not natural, and sometimes, if a staircase isn't designed well it can actually be a little destabilizing, we have to sort of adjust our feet. An escalator does this for us. It automates that machinal behavior and when it doesn't work (which is rather often the case) then we have to reactivate that behavior. for a similarly although much more complex process is affecting our ability of orienting ourselves: we don't use maps anymore, which required some level of sort of lucid behavior

and attention, thinking about directing yourself in space, etc. we use Google Maps.

My hypothesis is that those kind of functions and applications, services can be considered as mental penumbra; those machinal behaviors that we now share with technical systems. And so technical systems for a long time were thought of exactly—and this is where comes back to the fetish of automation in the beginning—as being purely automatic. If a machine was replacing us it's because it was doing something automatic, purely repetitive; we didn't have to think about at all. What we're seeing today are technical systems that are not only doing things that are automatic but are also doing things that are machinal, that are developed through imitation, trial and error and communication, that have to be taught, that aren't natural, that aren't instinctive,

Of course, walking up stairs or reading a map don't require electricity or complex technical operations, like an escalator or Google Maps. Which then points to the labor and infrastructure of maintenance behind these technical systems. Things we don't see but that are actually very real and very active. Teams of people fixing metros, constantly maintaining the infrastructure that's working in our collective subconscious, the invisible movements of the social division of labor. We get a sense of this when things break down on a large scale. For example, when Meta servers, a few months ago, crashed. Millions of people around the world couldn't access Facebook, Instagram, WhatsApp for hours. And it was as if people didn't know what to do with themselves anymore, in a very embodied sense. They couldn't do certain things because they had forgotten maybe how to do them without these services. So the services work for us on a daily level on this machinal, behavioral level.

Moments of power outage, for example, when they're when there's no electricity and you have to get candles out: you realize that actually your freezer or fridge doesn't work. So we've always shared bits of our behavior with machines. We don't have to salt our meat in the same way; we've lost a whole way of storing and cooking food because we have fridges today, and we like that. But when there's a power outage you're like "oh no, how do I preserve my food now?" And so it's true that outages or even a strike, when there's no more public transport, people have to organize differently and develop new behaviors.

So I think they're interesting moments where this big technical systems are actually becoming

so big that they're becoming also very vulnerable, fragile. I don't see that as a problem or a disaster. It's really a good thing that these technical systems are vulnerable, that they're weaker than we thought they were, as long as we develop a technical culture, institutions, and practices where we can share knowledge and care for the systems with which we already share so much. And perhaps also better determine which systems we do not want.

Further readings

- Bartholeyns, Gil and Charpy, Manuel 2021. *L'étrange et folle aventure du grille-pain, de la machine à coudre et des gens qui s'en servent*. Paris: Premier Parallèle. (OCLC: 1285483877)
- Bratton, Benjamin H. 2016. *The stack: on software and sovereignty*. Cambridge, Massachusetts: The MIT Press. Retrieved from <https://doi.org/10.7551/mitpress/9780262029575.001.0001> (OCLC: 960929264)
- Crawford, Kate 2021. *Atlas of AI: power, politics, and the planetary costs of artificial intelligence*. New Haven: Yale University Press. (OCLC: 1111967630)
- Denis, Jérôme and Pontille, David 2022. *Le soin des choses: politiques de la maintenance*. Paris: la Découverte. (OCLC: 1350520854)
- Dumouchel, Paul and Damiano, Luisa 2016. *Vivre avec les robots: essai sur l'empathie artificielle*. Paris: Éditions du Seuil. Retrieved from <http://catalogue.bnf.fr/ark:/12148/cb44515214j> (OCLC: 944248480)
- Reigeluth, Tyler 2023, 5. *L'intelligence des villes. Critique d'une transparence sans fin..* Météores.
- Simondon, Gilbert 2013. *Sur la technique, 1953-1983*. Paris: Presses Universitaires de France. Retrieved from <http://catalogue.bnf.fr/ark:/12148/cb437520160> (OCLC: 870997297)

