RESEARCH ARTICLE

Time Machines: Artificial Intelligence, Process, and Narrative

Mark Coeckelbergh[1] 🔾

Abstract

While today there is much discussion about the ethics of artificial intelligence (AI), less work has been done on the philosophical nature of AI. Drawing on Bergson and Ricoeur, this paper proposes to use the concepts of time, process, and narrative to conceptualize AI and its normatively relevant impact on human lives and society. Distinguishing between a number of different ways in which AI and time are related, the paper explores what it means to understand AI as narrative, as process, or as the emergent outcome of processes and narratives. It pays particular attention to what it calls the "narrator" and "time machine" roles of AI and the normative implications of these roles. It argues that AI processes and narratives shape our time and link past, present, and future in particular ways that are ethically and politically significant.

Keywords Artificial intelligence (AI) · Time · Process · Narrative · Process philosophy · Ricoeur · Bergson

1 Introduction

Artificial intelligence (AI) and data science, in particular machine learning and its contemporary applications that enable automated classification, prediction, decision-making, and manipulation in many domains of human activity and society, has sparked much controversy during the past decade: there are many concerns about the (potential) ethical and societal impact of AI. In philosophy and related academic fields, this has been accompanied by a wave of publications on the ethical and political aspects of AI (Bostrom, 2014; Boddington, 2017; Eubanks, 2018; Benjamin,

---

This article is part of the Topical Collection on *Philosophy of Technology and the French Thought*.

✉ Mark Coeckelbergh
mark.coeckelbergh@univie.ac.at

1 Department of Philosophy, University of Vienna, Universitätsstrasse 7 (NIG), 1180 Vienna, Austria

2019; Zuboff, 2019; Danaher, 2019; Couldry & Mejias, 2019; Coeckelbergh, 2020; Véliz, 2020; Berberich et al., 2020; Mohamed et al., 2020; Bartoletti, 2020; Crawford, 2021; Santoni de Sio & Mecacci, 2021; Cowls, 2021). However, less attention has been paid to the philosophical nature of AI. What are we talking about, if anything, and how can it best be conceptualized?

One way to conceptualize AI and data science and to understand how they impact our lives and societies, is to focus on time and related concepts from the philosophical tradition such as process and narrative. While the topic of time and technology has already received some discussion in philosophy of technology, for example in the Heideggerian tradition (Hui, 2012, 2017; Stiegler, 1998), analytic philosophy of robotics (Seibt, 2017), and earlier in the work of Virilio (1977) and Mumford (1934), and narrative approaches to technology have been proposed based on Ricoeur (Kaplan, 2006; Reijers & Coeckelbergh, 2020; Reijers et al., 2021), more work is needed on concrete contemporary technologies such as AI, and on the precise relations between the concepts of AI, time, process, and narrative.

This paper addresses this need by (1) exploring what it means to take a process approach to AI, by (2) introducing a distinction between 3 different ways in which AI and time are related, and by (3) showing how each of these conceptualizations have normative implications for how we evaluate the impact of AI on our lives and societies. One of the relations and conceptualizations I will propose concerns AI's role as a "time machine": I will argue that via classification, prediction, and recommendation, AI links past, present, and future in particular ways, which has normatively significant consequences. Moreover, throughout the paper I will explore what it means to see AI not as a thing or agent, but as process, narrative, and emergent *outcome* of processes and narratives. I will conclude that this process-oriented and hermeneutic approach has also implications for how we understand the human and for the way we (should) do ethics and politics of AI.

2 AI as Object Versus AI as Process and Narrative

We are used to think of technologies in terms of objects. When we think of technology, we imagine material objects such as a hammer or immaterial objects such as data and software. When we think of AI, then, we imagine a computer, a robot, software, a self-driving car, or other *things*. This way of seeing AI, which is in line with a long-standing tradition in Western metaphysics since Plato that sees the world as a collection of objects or substances, is reflected in philosophy of technology and has not changed in contemporary times. On the contrary, after the so-called empirical turn (Achterhuis 2001), philosophers of technology focus on artefacts, for example in the tradition of postphenomenology: what matters is what 'things' do (Verbeek, 2005). Information and informational technologies have also been conceptualized in terms of objects. For example, Floridi has argued that the world is the totality of informational objects dynamically interacting with each other (Floridi, 2014).

However, in Western metaphysics we also find another tradition, process philosophy, according to which the world is not a collection of objects but a process of becoming (rather than being). This tradition finds inspiration in Heraclitus' doctrine

of radical flux (consider the famous slogan *panta rhei*: everything flows), has been developed in German idealism (Hegel) and pragmatism (James, Dewey, Mead, Peirce), and is also to some extent present in Heidegger, but has found its most famous elaboration in the process philosophy of Bergson and Whitehead, which has then influenced contemporary philosophers such as Deleuze, Simondon, and Latour. The French philosopher Henri Bergson argued that individual intelligence emerged in a process of evolution that expresses a life force (*élan vital*) (Bergson, 1907). He used the term 'duration' (*durée*) to talk about time: already in his doctoral dissertation and in his debate with Einstein, he distinguished between time as we experience it (lived time) and the time of science, which is conceived of as discrete, spatial constructs (Bergson 1889). Bergson argued against what he saw as Kant's mistake to think time in terms of space, and emphasized intuition (Bergson, 1896) rather than formal knowledge and its conditions. However, duration is not just subjective, psychologically experienced time (as Einstein thought). Duration is something real, which can be experienced or transformed into something spatial. Thus, according to Bergson, there is not first objective time and then our experience of that objective time. Bergson attempted to go beyond the subject-object divide of Cartesian and Kantian philosophy: time cannot be isolated from our experience of it and, more generally, from 'the living structures in which it manifests itself' (Landeweerd, 2021, 26–28). What we call objective time is produced by our instruments: by technology. Physicists such as Einstein misleadingly make metaphysics out of this time-making. But the only metaphysics we need, Bergson argued, is one that recognizes duration and stresses emergence. The English philosopher Alfred North Whitehead, who thought at Harvard, used the term 'process': actual existence is a process of becoming (Whitehead, 1929). Whereas Western philosophy has traditionally privileged being over becoming, process philosophy reverses this. Moreover, like Bergson, Whitehead wanted to go beyond the subject-object divide: he sought to fuse the objective world of facts with the subjective world of values. In his process metaphysics, entities and experience are both part of becoming.

What would it mean to conceive of AI as process, lived experience, and becoming? This is difficult to imagine, since we are used to see and imagine AI as a some-thing, a substance. For example, we may observe the results of a statistical model (which we see as a thing), see a computer with AI software, or imagine a car that is driven by an AI system. This way of perceiving AI already bifurcates the world into a perceiver and some-thing that is perceived. We may also observe AI at different times. Time is then constructed in a spatial way: as a succession of discrete moments. At time t1 "the AI" does x, at time t2 "the AI" does y, and so on. This scientific way of understanding AI can be contrasted with our personal experience of the technology. For example, driving in a self-driving car — or rather *being* driven by an AI system — can be experienced as a flow, rather than a succession of distinct moments. Seen from a traditional modern view, these different concepts of time clash: there is a gap between the technology and the lifeworld, between objective and subjective knowledge. However, with process philosophy, we can radically question the metaphysical basis for such a gap, and see AI, *and* how we relate to AI and experience AI, as a process rather than an object, and more specifically a process that fuses "objective" and "subjective" elements,

science and lifeworld, scientific time and lived time. In the next section, I will further show what this means.

Moreover, if AI becomes part of the lifeworld at all, this is always mediated by our (human) interpretation and narration. This leads us to another interesting tradition in philosophy that is concerned with time and experience: hermeneutics, and in particular hermeneutics that focuses on time and narrative. Here Paul Ricoeur's narrative theory, which goes back to Aristotle, is most relevant. Ricoeur relates temporality to narrativity: our understanding of time and the world is mediated by language, in particular by narrative (Ricoeur, 1980, 1983). Ricoeur, too, is interested in bringing together different ways relating to time, but this time not process but *narrative* does the work: he argues that 'time becomes human to the extent that it is articulated through a narrative mode' (1983, 52). This takes the form of emplotment. Inspired by Aristotle's *Poetics*, he describes how the plot of a narrative configures characters, motivations, and events into a meaningful whole.

Recently, Ricoeur's theory of narrativity has inspired work in philosophy of technology. Kaplan (2006) already argued that, while Ricoeur's own view of *technology* is not so interesting since it belongs to a tradition that equates technology with domination, his work on *narrative* can nevertheless be used in philosophy of technology: technology, like a text, gains meaning from a background (use-context) and is open to multiple interpretations (49). Technologies also figure into our lives in different ways; we can tell stories about this (50). Emphasizing the more active narrating role of technology, Reijers and Coeckelbergh (2020) have used Ricoeur's narrative theory to show how technologies, similar to texts, have the capacity to tell stories and configure our lifeworld. Based on Ricoeur and in dialogue with the virtue ethics tradition, they have proposed a hermeneutic approach to ethics of technology that combines a focus on practices with narrativity theory: the way technology configures practices is linked to narratives in a community and to normative ideals.

What does this hermeneutic approach mean for AI? In what way could AI be like a text, or tell stories? And what does it mean when AI configures our lifeworld and is linked to broader narratives and norms?

One example of how to do this, which also shares the aim of trying to close the gap between technology and lifeworld, is offered by Keymolen (2021), who uses Ricoeur to analyse the Alpha Go AI as a story. By engaging with the story, she argues, we can interact with its power even before it is part of our everyday life (252). In her analysis of a documentary about Alpha Go, she shows how the game Go pre-structures the story of the AI and gives specific roles to the participants, how the story is embedded in a history of games between humans and technologies (259), and how 'tragic-like emplotment' plays a role, giving us an aesthetic experience (261–262). Explainability, then, is not just about opening the black box (giving technical insight into AI) but also about helping people to interpret and gain access to AI in a narrative, interpretative way. In other words, AI is not only a technology but also a story. This enables us to critically examine AI from a hermeneutic point of view and connect to (other) stories available in our culture, which give meaning to AI. It also helps us to analyse the structure of the stories (in this case a game structure).

Another example is the use of AI in contexts and practices of medical diagnosis. AI, by offering a diagnosis based on image recognition, can be understood as shaping the narrative of patients and doctors, giving them roles in a story about images and probabilities and influencing meaning making in this particular con-text. One could also understand this as a re-shaping of a medical *practice*, which is itself linked to narratives in a particular medical community or patient community, and which is subject to ethical and political norms and ideals (as for example materialized in medical ethics codes and laws).

Having explored some ways in which AI can be seen in process and narrative terms and having sketched the theoretical background of this paper, let me now propose a more general conceptual framework for thinking about the relations between AI and time, which combines some insights from process philosophy and narrative theory, and which also has normative implications.

3 Three Relations Between AI and Time, Described by Using the Concepts of Process and Narrative

Let me distinguish 3 relations between AI and time, which I name as follows:

1. The time of AI
2. AI in time
3. AI-time

3.1 The Time of AI: Narratives About AI

The first relation, titled "the time of AI", concerns the narratives we tell *about* AI. This can take place at a "macro" level. Consider for example the transhumanist accelerationist and Singularity narrative (Bostrom, 2014; Kurzweil, 2005; Moravec, 1988), according to which we are on the way towards a situation in which super-intelligence surpasses human intelligence, takes over control, and spreads into the universe, or the Marxist history of AI-capitalism, which predicts a future without humans (Dyer-Witheford et al., 2019). AI is then not just a technology but also a story: a story about civilization, about a particular society, about capitalism. AI is not only developed and used, but also narrated and interpreted. This can be done by putting AI in the light of these macro narratives. But narrating AI can also be done at the "micro" level: for *my* life. It can take the form of a particular story about me and AI. For example, someone could tell a personal story of how an AI did not recognize her, that this was really offensive to her, and that it shows how racist her society is — which in turn may lead to telling other stories, for example stories of discrimination in context of police intervention. In both cases, the meaning of AI takes shapes and evolves, against a background that helps to give it meaning. To pick up Wittgenstein's term Ricoeur draws attention to in his article on Wittgenstein and Husserl (Ricoeur, 2014): like a text, the narration, construction, and interpretation of AI take place against the background of a 'form of life', which helps to

give meaning to AI and which is in turn constituted by the specific meaning-giving activities of people — concerning AI and otherwise.

Again this means that AI is not just a "thing" (although it certainly has material aspects — see for example Crawford, 2021), but also a particular narrative or collection of narratives, which are linked to other narratives. The stories are not just "about" AI, as if AI were a fixed thing that is not influenced by the stories we tell about it. By telling stories about AI, humans also constitute what AI "is," or rather, what it becomes as a result of the stories. This is literally the case, in the sense that for example transhumanist narratives might influence the actual research and development of AI. But if we consider AI as use and practice, narratives about AI also shape how we interact with AI, how we think about AI, how we talk about and to AI, and so on. One could also say that there is no AI-in-itself (to use a Kantian phrase); what AI becomes is shaped by narratives: the narrative about the particular AI, but also "grand narratives" about the history and future of humanity.

Consider for example Harari's (2015) narrative that re-tells the history of humankind in a way that predicts a future in which humans will be obsolete when intelligent machines and "Dataism" takes over. Harari's narrative is a contribution to the mentioned transhumanist grand narrative about humans that enhance themselves and create intelligent machines, which leads eventually to a so-called Singularity and intelligence explosion: new, artificially intelligent entities take over and spread into the cosmos. No need for humans, with their limited intelligence and limited data processing capacities. Specific AI technologies and technological events (so-called "breakthroughs", for example, or demonstrations of AI's power as in the case of Alpha Go, GTP-3, or Wu Dao 2.0) are then seen in the light of this grand narrative: as steps on the way towards the Singularity. A narrative about a particular AI is understood as part of a larger narrative, which gives meaning to particular technological breakthroughs.

Note also that these narratives are structured in a particular way. The Singularity and intelligence acceleration and explosion narrative, for example, borrows its "grammar" from Moore's Law, which concerns the rate of growth of computer hardware and thereby speed: the observation that the number of transistors in an integrated circuit doubles every two years. The grand narrative of the future of humanity and the universe is structured by this: the larger narrative is itself based on a "smaller" narrative about hardware development and computer power, which is a very specific way of perceiving/constructing time and AI.

3.2 AI in Time: AI and Data Science as Process

The second relation, titled "AI in time", is again about AI as a process, rather than a thing, and refers to the use and development processes that take place in time, for example AI as data science process. "Time" here can refer to two different times in which AI takes place or develops: scientific-objective time and the time of the lifeworld, the lived time Bergson conceptualized as *durée*. However, in process, the two kinds of time merge: AI in time is then both measured/controlled and lived. It is a duration that is both experienced by humans (lived) *and* rendered "objective"

and produced by measurements, technologies, and management techniques. The best way to understand how this happens is to consider data science processes. These processes have various steps, such as data collection, data analysis, modelling, and so on. This way of perceiving/constructing the data science process belongs to what twentieth century philosophers would call "objective" time or modern-scientific time. It is about management and control. The steps divide up time in a way that renders it spatial. The different steps are different boxes, marking discrete chunks of time. But every step involves humans, who experience, act, and interpret. There is not only the time as shaped by the technological and scientific process; there is also human experience and human experience of time. In data science as a practice, both the technological-scientific process and the lived time are at work. Conceptually both kinds of times can and must be distinguished. But in process and in practice they combine.

What AI "is", then is this process, or even the outcome of the process. It is impossible to say what AI "is" a priori, before or outside the process. AI and the data science processes are connected. AI cannot be "lifted out" of time, and neither can it be disentangled from what humans do and experience. The process can be described in spatial terms (in terms of steps), but knowledge of the process is always at the same time lived. Moreover, at the same time AI also leads to the emergence of human subjects: the measurer and controller are an outcome of the measurement and control process. The data scientist is shaped by the data science process.

3.3 AI-Time: How AI Shapes Our Time and Functions as a Time Machine

The third relation, titled "AI-time", differs from the previous ones in at least two ways:

First, here AI is neither a thing or object, nor *just* (part of) a narrative or process, but becomes itself a meaning-maker, a writer of the story: AI becomes a co-narrator, rather than just an object or even actor in the story. In this role of narrator, AI shapes our time and is a "time machine", which links and shapes past, present, and future.

What does that mean? Let me unpack this "time machine" and more active "narrator" role. First, by making *classifications* based on historical data, AI processes may fix us in the past, thus shaping particular presents and futures. For example, if historical data from job interviews are used to train a hiring algorithm, then past ways of thinking — including potential bias — will shape present hiring and thus the future of the company and the story of the people who are (not) hired. Second, by means of *prediction*, which then influences human action, AI processes shape the present and future. For example, if AI predicts that there will be more crime in a particular area, then police forces may focus their activities there and prevent more crimes in that area, which changes the present and future. AI then creates a self-defeating prophecy. Third, by making decisions and by manipulating people, AI shapes the present and future. As a decision-maker, AI becomes a character in the story. As a manipulator, it becomes a co-narrator of the story. For example, if it takes on the role of a judge deciding about parole, it is an actor in the story (an automated judge) and it co-writes the story of a particular prisoner and, in the end, it shapes the

history of the decisions of that particular court. This in turn feeds into future uses of AI. And if AI is used for manipulating people, for example for nudging them to buy certain products by making recommendations based on their statistical profile (consider for instance Amazon), then it changes the story of that particular consumer.

AI does not play this role as a "thing", but rather as (1) a process that is structured in particular ways and as (2) a narrator of human lives. While AI is not a narrator in the same way as human beings can be narrators, since AI is not an intentional and sensemaking being (and indeed no being at all but rather a process), AI co-shapes our narratives. Given the inherent structure and functions of AI and data science as *processes* and *narrators*, AI works as a time machine. If we only use terms such as "artefacts", "objects", and "things", we cannot conceptualize this: we need a process-oriented and narrative approach, which relates AI to humans and their activities, structures, and culture.

Second, however, if we further radicalize this approach in a process philosophy direction, there is no longer an opposition between AI and humans as fixed relata *in* processes and stories; instead, they emerge from the process itself. I already suggested this with regard to the "AI in time" relation, but here this becomes even clearer. First there is the story, the process, the relation. We do not start with fixed entities; what we call "AI" and "humans" emerge from the process, they become. For example, in the manipulation case what "AI" is, becomes clear in and through the data science process that leads to the manipulation; it cannot be defined separate from that process and is rather the result than an ingredient or tool. It is the cake, not (just) the mixer. Similarly, the human in this process and story is not fixed from the beginning but becomes what she is through the process and by having received a role in and through the story: she starts with the idea that she is an autonomous individual, perhaps, but is then made into a manipulated consumer in and through the process. AI as (part of a) process and narrative (for example a marketing process and a capitalist narrative) gives her that role. And if she resists, protests, and so on, then she starts a new process and story, which connects to the existing story and may or may not lead to a different outcome. (I will soon look further into the normative significance of this analysis.)

Interestingly, this means that we are not fully in control of AI: not in the sense that AI does things without our intervention (consider again the self-driving car), but in the sense that it can get *hermeneutically* out of control and perhaps always is out of our full control in that sense. We do not fully control the meanings and roles that are the outcome of AI and data science processes. The developers may have one interpretation of what their AI "is" and "does"; but what it *becomes* can be very different since other interpretations are possible and since the outcome of a process is not always fully predictable.

Is what AI does here and means here unique? Is it different from other technologies? And keeping in mind the Ricoeur-inspired hermeneutic approach: is it different from text?

Yes and no: One the one hand, these relations and roles of AI are not so different from what text does. Text is also a technology we can talk about, a process, and meaning-maker. It also has emergent properties, and we are also not necessarily in control of the meanings and roles that emerge. The author has long been

proclaimed death (Barthes, 1967); the author does not fully control the meaning of the text. This seems also true for the developer, whose intentions may clash with what (end-)users do with the program. And as we know from the tradition of thinking about writing technology from Plato to Stiegler, technologies also constitute a kind of memory. In the *Phaedrus*, Plato already worried that people would cease to exercise memory because they would rely on writing. Printed text can be seen as an extended memory (Ong, 2012). Like text, AI and data science processes fix knowledge of the past. Once it is on the page (text) or in the dataset and processed by the algorithm (AI), there is no real-time change anymore. Just as in text we might get bewitched by the thoughts and stories of the past, data science processes may prevent social change by perpetuating biases of the past. At the same time, however, there is no determinism. We can offer different interpretations of the text and we can change the algorithm, the data, and (in principle at least) human behaviour.

On the other hand, there are at least three differences with AI. First, AI produces a different kind of knowledge: not text (say linguistic knowledge) but numbers, in particular statistical knowledge, for example probabilities, correlations, etc. AI and data science processes are therefore not an exteriorization of human memory, as Stiegler (1998) saw it, but amount to a different kind of memory altogether. AI and data science processes produce their own kind of knowledge, which is then memorized in technical ways (databases, models). Whereas the Platonic model of writing presupposes some pre-existing memory in the human, which is then exteriorized through writing and materialized on paper, AI and data science transform human thought and experience into data, and produce statistical knowledge about these data, which the humans involved do not already have and (especially in the case of big data and complex models) *cannot* have or produce. AI thus creates its own "memories", which may be quite different from the content of human memory, which is based on human experience and not on data.

While it is well-known in the phenomenological tradition (Husserl, 1973; Merleau-Ponty, 2002) that there is 'sedimentation' in the sense that past experiences influence present experiences of the same phenomenon, in the case of AI processes both the past and the present reliance on that past take a very specific shape, which is not about human experience but about the production and use of statistical knowledge. While humans are part of the process of AI and data science, the transfer from past to present itself is done without human experience intervening. There is no 'sedimentation' in the traditional, phenomenological sense of the word: instead, there is calculation. At best, there is interpretation by humans afterwards (which then can be the object of sedimentation).

There may be 'sedimentation' of the *technology* in the sense that the use of technology itself may recede into the background (Rosenberger and Verbeek, 2017; Lewis, 2021). This phenomenon is well-known since Heidegger and Merleau-Ponty (and later Dreyfus). In the case of AI, we may for example use a search algorithm but not be aware that it relies on AI. (It is not sure, however, that this can be described in terms of sedimentation, since what happens there is different from the incorporation and creation of embodied knowledge that for example Merleau-Ponty and Dreyfus describe, and there is no gradual receding into the background since on the part of the

user, the technology may never have been in the foreground in the first place: we were never aware of it, it was hidden.)

But here I consider a different phenomenon, which has to do with transfer of knowledge from past to present. In the AI and data science process, there is no sedimentation in the sense of human experience that provides the basis of further experience. The knowledge produced and used by AI is not directly based on human experience. It is only very indirectly, via data, that human experience plays a role. Text, writing, and narrative seem to offer a more direct access to human experience, albeit never totally direct and always mediated. But this way of putting it is also not right: it misleads us into thinking that there is always a pure human experience in the first place. Instead, the technology of writing co-shapes the experience; it does not fully pre-exist as a fixed kind of thing that is pure and untouched by technology and/or its linguistic and narrative expression. Similarly, both the AI knowledge and the interpretations by humans do not pre-exist but *become* in and through the AI and data science processes in their context of application. Just as text is not just the mirror of knowledge that is pre-existing — one could say that the meaning and knowledge *becomes* in the process of writing — AI knowledge becomes during the process. Yet this becoming is not as open as what happens when humans write. During the data process, considered only in its technical aspects, there is no sedimentation in a phenomenological sense since this is not about human experience in the first place; at most, there is a technical process of memorization. The algorithm and the model themselves do not involve interpretation, although the result does. This leads me to the next point.

Second, we are used to see writing and text as something that requires interpretation and hermeneutic dialogue and communication: between reader and text, between readers. The meaning of the text is not limited what the author intended, and the writing itself was already a hermeneutic process: the meaning of the written text evolved in "dialogue" with other texts and meanings available in the language and culture. AI, by contrast, is seen as an instrument, and tool, a thing, that is hermeneutically neutral: we (the developer and the user) are the ones who give meaning; AI and other technologies are supposed not to be "hermeneutically active" or "hermeneutically creative" themselves. But as I have argued, this assumption is mistaken. AI technologies, like digital technologies in general (Romele, 2020), are interwoven with meanings and also change these meanings. It's use and meaning go beyond what the developers intended. And as we have seen, AI can be understood in terms stories: stories that help us to give meaning and stories that are always related to other stories. Therefore, it is necessary to develop a hermeneutics of AI. If we see AI as a "thing" and "technology" only, we fail to do this; we believe that hermeneutics is a matter of text only. But after the conceptual work done in this paper and the sources it refers to,[1] the road to a hermeneutics of AI is open.

---

[1] Note that there are also other sources, for example in the work of Ihde (1998) we find the concept of 'expanding hermeneutics'. However, in elaborating this, Ihde does not *directly* draw upon the hermeneutical tradition (let alone process philosophy), whereas I engage directly with Ricoeur.

Given the kind of knowledge presented by AI, however, such a hermeneutics is quite a challenge: how to get from statistical knowledge to knowledge that can be used by humans? If we ask this question in the abstract, it is quite a mystery how this happens. If we look at AI and data science practice, however, we see that this translation is already done by humans: humans who know data science and humans who interpret and can relate the knowledge presented by AI to their lifeworlds and narratives. By seeing AI as a thing or by only considering knowledge creation by AI in the abstract (as I did previously when I talked about technical processes that move from past to present), these humans and their interpretation were removed from view. But by seeing AI and data science as processes that involve both technologies and humans, we gain conceptual space to talk about this challenge. For example, we can look at how humans take decisions about data(sets) and models, and we can observe and discuss how statistical knowledge is and should be interpreted, and what role we want to give to that kind of knowledge in our lives and our society. This is then not a question of "science communication" — if that means that scientists explain to the wider public — but of common *interpretation*. I propose that we develop a political hermeneutics of AI, and, more generally, of science and technology, which describes and evaluates political-technological processes of common meaning making.

Third, as some of my examples already suggested, AI hermeneutics differs from classic text hermeneutics in its basic orientation to the past. Paradoxically, while classic hermeneutics often deals with ancient texts, it is fundamentally oriented towards the present and future. The focus of classical hermeneutics is on rendering the old texts relevant to contemporary readers and indeed to contemporary times. The dialogue between reader/interpreter and text is oriented towards the future: the text should become effective for the reader. It should be used to shape the present and future of our lives and communities. AI, by contrast, does not only deny that dialogue, but also risks to close of the future. As it classifies and predicts on the basis of past data, it risks to treat people in ways that nail them to their past and to what cannot be changed (any longer), for example their ethnic origin or their past desires and past behaviour.

This discussion leads us to the normative implications of AI as time machine and meaning machine.

4 Normative Implications of AI as Time Machine and Meaning Machine

Let me show how this analysis of relations between AI and time in terms of process and narrative also has normative implications.

First, the *narratives* to which we link AI are not normatively neutral, but have their own built-in normative recommendations and demands. For example, if I entertain a humanist narrative about modern technology as a threat to human culture, then I will see AI as a problem, perhaps even something to resist and fight against. If, on the other hand, I embrace a transhumanist narrative, then I will see current AI as a step in the welcome development of superintelligence and the enhancement or

superseding of humans in the light of this grand history of the cosmos. If we see AI as a thing or a technology only, we miss these normative-narrative dimensions. We then believe that narratives, morality, and politics belong to a different world, not the world of science and technology. This excludes the possibility to critically analyse and discuss these narratives in a productive way.

Second, AI and data science *processes* with their different steps and human interventions and experiences are also not normatively neutral. For example, at various stages of the process, bias may slip in, both as a result of automated data collection (e.g., biased data harvested from the internet, which contains text corpora that are biased) and as a result of human action (the team of data scientists may be biased and take biased development decisions). Instead of seeing AI as a biased "thing" or "agent", we can now talk about the biased nature of the process as a whole, and present a more precise analysis of bias at different steps and *times* in the process. This temporal aspect is often missing in ethical analysis of AI. Moreover, there is not just the "objective" time and the steps; these times and steps are themselves created by humans, and hence can be morally and politically evaluated. Nothing data scientists do is "natural" or "set in stone". "AI-time" is human work, even if humans do not have full control over its outcomes and meanings.

Third, if and to the extent that AI shapes our time, this is also normatively relevant. Consider the various *time machine* functions. If AI fixes us in the past and prevents social change, for example by reproducing historical bias, then this is normatively significant, including politically significant. In order to do something about this, we may then "resist" AI or call for *different* time machines, which do not incur bias or even repair it. If, as I argued in the previous section, the temporality of AI is oriented towards the past, then we need to try to re-orient it towards the future. We need to open up the possibility for dialogue and interpretation. It would also be good if the stories of those impacted by AI (or those who may potentially be impacted by AI) could somehow become part of AI and data science processes, rather than — in the best case — being presented afterwards in the form of criticism in the (social) media or taking the form of exercising a legal right of reply, which usually comes too late. This is both a technological and a normative project, and it is a matter of time. We need to figure out how to *make time* in a different way: how to let in the future and make time for social change, how to make time for interpretation and judgment, and how to make time for people and their stories.

The narrative role of AI has similar normative implications. If, through prediction and recommendation, AI shapes our choices and actions — and thus shapes our narrative and our future — then this is also normatively important. Consider the predictive policing example: if less people commit a crime in a particular neighbourhood as a result of actions taken recommended by AI, then this is ethically and politically good (assumption: reducing crime is good). At the same time, if people in that same neighbourhood feel targeted as a result of the police actions, this is ethically and politically problematic: people may invoke a narrative about historical discrimination with regard to people living in that area, people with a particular ethnic background, for example. Prediction by means of AI, understood as narration and having a "time machine" function, is thus normatively significant and may invoke entire political discourses and narratives about poverty, race, discrimination,

(neo)colonialism, and so on. Personal and collective experiences from the past and visions of the future, developed in a specific time and geo-cultural context, will play a role in these discourses and narratives. In this sense, what AI "is" and "does" in for example the United States may be different than in a European country, since there are not only different data but also different interpretations involving different histories and different visions of the future. AI tells different stories and, as a time machine and "meaning machine", moves between different times and brings back different pasts (e.g., the time of slavery in the USA versus the time of colonialization of a particular African country such as Leopold's Congo). For example, Ruha Benjamin's book (2019) stresses race and talks about white supremacy and abolitionism because in the USA this is part of its specific history and present; one could say that the book shows that AI (re-)narrates a story about race, social inequity, and white supremacy. Of course what she says about AI is partly universally applicable, but the point is that AI works here as a time machine in a specific context and history. Understood in this way, AI is not just an instrument for data analysis but becomes part of the meaning-making tools available to deal with the past and (re-)shape a present and future. AI is then a tool for a particular society's political hermeneutics. We can and should criticize the stories written with and by AI, and create new, better narratives.

However, again we should avoid talking about this as "AI" having particular normatively relevant consequence and hermeneutically relevant roles, as if it were an agent or thing on its own. It only fulfils its role as "co-narrator" and "time machine" as part of mixed human/non-human processes and narratives. What AI "is" or, rather, becomes, emerges from those processes and narratives. For example, if in a particular present situation and process bias is created or people are exploited, AI may *emerge* as the face of political discrimination and neo-colonialism as a result of the process and in a particular narrative and context. Similarly, human subjects emerge from these processes and narratives. This may involve all kinds of narratives, including myths and hero stories. For example, in the transhumanist narrative, an entrepreneur who develops general AI may emerge as a hero who shapes the future of humanity (or helps to bring about a future without humans), and in a particular data science process context, an ethicist who stands up against big tech may emerge as a rebel. Humans whose data are collected and sold or humans who do dirty work for producing AI-powered devices may emerge as victims of the data industry and data capitalism. And a particular group or community in a city may emerge as victims of predictive policing. All these roles emerge from processes and stories in which both humans and AI act, figure, and co-author.

This discourse about humans and technology in terms of emergence is not the same as, but fits with, Bergson's view of technology as a manifestation of the vital impulse and as individuation (Bergson, 1896) — which later influenced other French philosophers such as Simondon and Deleuze. Simondon (1989, 2005) saw both individuals and technology as outcomes of processes of individuation.[2] To

---

[2] Note that there is an interesting connection between Simondon, cybernetics, and the philosophy of information (Iliadis, 2013). However, I will not further discuss this here.

see technology as part of evolution and as the outcome of a process of emergence offers a different conception of technology than the usual instrumental one. As Landeweerd (2021) puts it:

> The advantage of such a reading of technology is that it circumvents the mistake of seeing humanity as the sole author of technology, as captain at the help of the ship of Progress, navigating through the unchartered seas of innovation, able to amend processes of diversion generated by technology. (Landeweerd, 2021, 143)

At the same time, humans are an important part of the processes and emergences. There is no technological determinism. We co-author the narratives and participate in the processes. And as suggested earlier when using Bergson, human experience (subject) cannot be disentangled from technology (object). In this sense, too, humans are part of the process and part of the story. We could call this "lived technology". AI, too, is a "lived technology" in this sense: there is no such thing as an AI-in-itself, AI is always experienced-AI. It is also always interpreted-AI.

For doing ethics and politics of AI, then, this process-oriented and narrative approach means that we should not only talk about "what humans do" with AI or not even at "what AI does" to humans, but also critically examine the process and narratives that lead to such constructions and emergences. Once we conceptualize how humans and AI relate as a process of becoming and as a matter of narrative meaning, we can look below the surface of the usual technical and normative discourses and try to understand and evaluate the processes, the narratives, and the structures of these processes and narratives from which humans and technologies emerge. Process philosophy and narrative theory give us conceptual tools to do this. The examples in the previous pages show that AI is very much entangled with small and big narratives, that technology and culture are intertwined in processes and stories.

5 Conclusion: Towards a Moral and Political Hermeneutics of AI

Inspired by process philosophy and narrative theory, I have explored what it would mean to conceptualize AI in terms of process, narrative, or an outcome of process and narrative, and I have analysed a number of ways in which AI relates to time. It turned out that once we *de-imagine* AI as a thing and dig into some important ways in which it relates to time, we can open up new avenues for understanding and evaluating AI and data science. The conceptual work presented here presents a process view of AI and offers elements for what we could call "a moral and political hermeneutics of AI". Beyond that, the proposed process-oriented and narrative approach may also be helpful to analyse other contemporary technologies — although here I limit the scope to AI.

Importantly, however, this paper was not only about technology but also about humans: the proposed approach removes humans from their supposedly central position as sole narrators of their stories (technology also narrates) and conceptualizes their role and subjectivity as emergent rather than pre-existing (humans as the outcome of processes, including technological processes). From a traditional humanist

point of view, this may be regretted, challenged, and resisted. From another, perhaps posthumanist point of view, it may be seen as an opportunity to critically analyse the role technology plays in narrating and making us and our societies. Nothing said here implies that one should uncritically accept the current AI-driven processes that construct us or embrace the stories that are told about us and AI. On the contrary, the conceptual framework offered can be interpreted as a stimulus to both critical analysis and change. And such a change is possible: luckily, we are co-makers of processes and co-narrators of our narratives. We can co-create new, better processes and stories. In the spirit of both Bergson and Ricoeur, one could say: we have the freedom to imagine and re-imagine, form and transform. Given this freedom and agency, we humans also have a responsibility for those processes and narratives. I propose that we exercise that responsibility, with regard to AI and otherwise.

References

Achterhuis, H. (2001). *American philosophy of technology: The empirical turn*. Indiana University Press.

Barthes, R. (1967). The death of the author. Trans. S. Heath. In Barthes, R., Image, music, text: Essays selected and translated by Stephen Heath, pp. 142–148. London: Fontana, 1977.

Bartoletti, I. (2020). *An artificial revolution: On power, politics and AI*. The Indigo.

Benjamin, R. (2019). *Race after technology: Abolitionist tools for the New Jim Code*. Polity Press.

Berberich, N., Nishida, T., & Suzuki, S. (2020). Harmonizing artificial intelligence for social good. *Philosophy & Technology, 33*, 613–638.

Bergson, H. (1889). *Time and free will*. Dover Publications, 2001.

Bergson, H. (1896). Matter and memory. Trans. N.M. Paul and W.S. Palmer. Zone Books.

Bergson, H. (1907). Creative evolution. Random House Modern Library, 1944.

Boddington, P. (2017). *Towards a code of ethics for artificial intelligence*. Springer.

Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.

Coeckelbergh, M. (2020). *AI ethics*. MIT Press.

Couldry, N., & Mejias, U. A. (2019). *The costs of connection: How data is colonizing human life and appropriating it for capitalism*. Stanford University Press.

Cowls, J. (2021). 'AI for social good': Whose good and who's good? *Introduction to the Special Issue on Artificial Intelligence for Social Good'. Philosophy & Technology.* https://doi.org/10.1007/s13347-021-00466-3

Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

Danaher, J. (2019). *Automation and utopia: Human flourishing in a world without work*. Harvard University Press.

Dyer-Witheford, N., Kjøsen, M., & Steinhoff, J. (2019). *Inhuman power: Artificial intelligence and the future of capitalism*. Pluto Press.

Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.

Floridi, L. (2014). 'Informational realism'. Available at SSRN: https://doi.org/10.2139/ssrn.3839564.

Harari, Y. (2015). *Homo Deus*. Harvill Secker.

Hui, Y. (2012). What is a digital object? *Metaphilosophy, 43*(4), 380–395.

Hui, Y. (2017). Rhythm and technics. *Research in Phenomenology, 47*, 60–84.

Husserl, E. (1973). Experience and judgment. Trans. J.S. Churchill and K. Ameriks. Evanston: Northwestern University Press.

Ihde, D. (1998). *Expanding hermeneutics: Visualism in science*. Northwestern University Press.

Iliadis, A. (2013). 'Informational ontology: The meaning of Gilbert Simondon's concept of individuation.' Communication and New Materialism 2(1). Available at https://scholarworks.umass.edu/cpo/vol2/iss1/5. Accessed 21 June 2021.

Kaplan, D. (2006). Paul Ricoeur and the philosophy of technology. *Journal of French Philosophy, 16*(1/2), 42–56.

Keymolen, E. (2021). 'Narrating artificial intelligence: The story of AlphaGo'. In Reijers, W., Romele, A., and Coeckelbergh, M. (2021). Interpreting technology: Ricoeur on questions concerning ethics and philosophy of technology, pp. 249–270. Rowman & Littlefield.

Kurzweil, R. (2005). *The singularity is near*. Penguin Books.

Landeweerd, L. (2021). *Time, life, and memory: Bergson and contemporary science*. Springer.

Lewis, R. S. (2021). *Technology, media literacy, and the human subject*. Open Book Publishers.

Merleau-Ponty, M. (2002). Phenomenology of perception. Trans. C. Smith. Routledge.

Mohamed, S., Png, M. T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology, 33*, 659–684.

Moravec, H. (1988). *Mind children: The future of robot and human intelligence*. Harvard University Press.

Mumford, L. (1934). Technics and civilization. The University of Chicago Press, 2010.

Ong, W. (2012). *Orality and literacy*. Routledge.

Reijers, W., & Coeckelbergh, M. (2020). *Narrative and technology ethics*. Palgrave.

Reijers, W., Romele, A., & Coeckelbergh, M. (2021). *Interpreting technology: Ricoeur on questions concerning ethics and philosophy of technology*. Rowman & Littlefield.

Ricoeur, P. (2014). The Later Wittgenstein and the Later Husserl on language. *Études Ricoeuriennes / Ricoeur Studies, 5*(1), 28–48.

Ricoeur, P. (1980). 'Narrative time.' Critical Inquiry, 7(1. On Narrative), 169–190.

Ricoeur, P. (1983). Time and narrative—Volume 1. In K. McLaughlin & D. Pellauer (Eds.). The University of Chicago.

Romele, A. (2020). *Digital hermeneutics*. Routledge.

Rosenberger, R., & Verbeek, P.-P. (Eds.). (2017). *Postphenomenological investigations*. Lexington Books.

Santoni de Sio, F., & Mecacci, G. (2021). Four responsibility gaps with artificial intelligence: Why they matter and how to address them. *Philosophy and Technology*. https://doi.org/10.1007/s13347-021-00450-x

Seibt, J. (2017). Towards an ontology of simulated social interaction. In R. Hakli & J. Seibt (Eds.), *Sociality and normativity for robots* (pp. 11–37). Springer.

Simondon, G. (1989). *Du mode d'existence des objets techniques*. Aubier.

Time and narrative—VolumeSimondon, G. 2005. *L'Individuation à la lumière des notions de forme et d'information*.

Stiegler, B. (1998). *Technics and time, 1: The fault of Epimetheus. Translated by R. Beardsworth and G. Collins*. Stanford University Press.

Véliz, C. (2020). *Privacy is power: Why and how you should take back control of your data*. Penguin/Bantam Press.

Verbeek, P.-P. (2005). *What things do: Philosophical reflections on technology, agency and design*. The Pennsylvania State University Press.

Virilio, P. (1977). *Speed and politics* (p. 2006). MIT Press.

Whitehead, A.N. (1929). *Process and reality*. The Free Press, 1978.

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile Books.