It's Time to Think About Time: Artificial Intelligence, the Speed of Organisational Culture, and the Social Acceleration of Knowledge Work

Advait Sarkar
University of Cambridge and University College London, United Kingdom
advait.sarkar@cl.cam.ac.uk

Not only is Artificial Intelligence changing the pace of work and culture, but also the pace of change itself. To the people living through it, this can feel like whiplash. To describe that change, we typically reach for languages of technological progress, economic growth, labour process, or occupational pressures. This paper introduces a different language from the philosophies of time, speed, and social acceleration. For more than a century, philosophers of speed have asked what speed does to culture, to power, and to the self, but their thinking has barely been brought to bear on the changes we are experiencing due to generative AI. We start from the position that AI can be read as a *chronotechnology*: a time machine that closes the gap between wanting a piece of work and having it, and that produces many futures at once. We then ask how AI acceleration reshapes the knowledge worker's relationship to the past, the present, the future, to power, and to risk. Our aim is to open the conversation between the discourse on AI and the philosophy of time, a conversation that we believe is currently missing, yet important and productive.

dromology, philosophy of speed, philosophy of technology, large language models, temporality, chronopolitics, time-space compression, desynchronisation, resonance, digital capitalism, future of work, temporal design

1. INTRODUCTION

I recently encountered an experience of dislocation likely to be familiar to many international travellers in the 21st century. I spent a late morning in Tarragona, Spain, drinking vermouth under a Mediterranean sky at twenty-one degrees Celsius. The same evening, I was walking through London in seven degrees of grey drizzle. It felt as though my body had arrived before my mind.

This is the whiplash of twenty-first century travel. The effect is a mild vertigo, a sense that space has been folded up and put away. The sensation is not new; it is the "annihilation of space by time" that writers since Marx have treated as the signature of accelerating modernity (Harvey 1989).

Something similar is now happening to knowledge work. With each advance in Artificial Intelligence (AI), the gap between imagining a piece of work and obtaining the finished artefact shrinks. Consider "project" to mean a scoped deliverable that today typically takes several months of full time work. In the not distant future, a knowledge worker might begin the day working on one project, which might be completed by mid-afternoon, and the evening project is something else entirely.

AI is also blurring and extending conventional roles, enabling knowledge workers to perform tasks traditionally outside their disciplinary expertise (Ulloa et al. 2026; Sarkar and Drosos 2025). So a software engineer might begin the day drafting a legal argument, and end it by performing accountancy.

Consider how this rapid transitioning across projects and roles might *feel* to the knowledge worker of the near future. Perhaps a kind of occupational whiplash: the same vertigo as induced in the modern jet traveller, now felt at a desk. The knowledge worker moves nearly instantaneously between states that used to sit far apart, and the work arrives before we have caught up with it.

Let us start with a simple observation: AI speeds up work, and with it life and culture. Many people who use these tools can sense this instinctively. Our question is whether we can put some philosophical meat, so to speak, on these instinctual bones.

I am, obviously, not the first to notice that technology speeds up the world. For more than a century, philosophers, sociologists, media theorists, and cultural critics have asked what speed does to human experience, and a rich body of work on

time, social acceleration, and the speed of modernity already exists. This is briefly reviewed in Appendix A.

Research discourse in human-AI interaction has not yet engaged these ideas. This paper is an attempt at an introduction. We offer a lens, AI as a time technology, a *chronotechnology*, and a set of questions we think are worth taking seriously. These include: how does AI affect the knowledge worker's perception of time? How does it affect their relationship to the past, present, and future? Does the acceleration of work redistribute power? What are the risks of speeding up work? How can we adapt to, and where appropriate, resist, the acceleration of work? These are simple questions that fall out of even the most superficial readings of the philosophy of speed, but are not being asked about the accelerating effects of AI. We believe that it is productive to do so.

Let us start with a simple provocation: AI is best understood as a *chronotechnology*, that is, a technology of time.

2. ARTIFICIAL INTELLIGENCE, TIME, SPEED, AND ACCELERATION

2.1. AI as a Time Machine

Ask a generative model for a report, an analysis or a working prototype, and what comes back is the finished thing: the artefact that would once have arrived next week, arriving now. In this sense AI is a kind of time machine. It shortens the distance between wanting and having, delivering the instant, on-demand present that Tomlinson calls immediacy (Tomlinson 2007), and which David Harvey wrote about as time-and-space compression (Harvey 1989).

But generative AI in knowledge work can have an effect rather different from simply accelerating a single timeline. Ask it once, and one future comes back. Ask it a hundred times and a hundred futures come back side by side. So AI is more than a time machine: it (ostensibly) frees us to explore, and even to inhabit, several possible futures at once. The time it opens up is less a line than a garden of forking paths: a private "multiverse" of future work.

This might seem liberating, but let us make two observations that serve as counterpoints to the assertions that AI is a time machine that frees us to explore. First, a technology that is time-saving in theory might be time-destroying in practice. The history of computing is littered with tools that promised speed and delivered its opposite – the "computer productivity paradox" (Landauer 1995). We will return to the possibilities of paradoxical slowdown in Section 2.6.

Second, going faster is not the same as being free. Byung-Chul Han argues that ceaseless activity, always busy, always doing, is the opposite of freedom. We have moved, on his account, from Foucault's *disciplinary* society of outside prohibition to an *achievement* society in which we drive ourselves: no master need stand over us with a whip, because we exploit ourselves, freely, to the point of burnout (Han 2015).

Notice, too, how trenchant is the tendency of the field to narrate itself in the language of speed. Progress is a race: higher benchmark scores, more compute, more investment, more parameters, more energy, bigger datasets, higher adoption, more tokens per second, longer context windows. Speed and acceleration are core idioms through which the technology understands itself. If that is the case, then AI must reshape how the knowledge worker stands in relation to time. In the following sections, we take that relationship in turn: the past, the present and the future; then the power that speed confers; and finally the accident it invents.

2.2. The Past: the Devaluation of Experience

As work speeds up, the useful life of what we currently know shortens. Hermann Lübbe called this the "contraction of the present": the shrinking window within which past experience remains a reliable guide for action (Lübbe 2009). He notes that Machiavelli could still hold up ancient Rome as a working manual for statecraft in Renaissance Florence, because the world changed slowly enough that the distant past stayed serviceable. Similarly, Henry Adams concluded that each generation would find it harder to instruct the next, because the world it was preparing them for had already moved on: *"The attempt of the American of 1900 to educate the American of 2000 must be even blinder than that of the Congressman of 1800 [to educate those in 1900]"* (Adams 1918). A workplace skill learned ten years ago might seem laughably quaint today, and a grandparent's career can read like a dispatch from a foreign country.

For knowledge work the implication is obvious. If AI makes the work move faster, the half-life of hard-won context falls with it: the playbook, the institutional memory, the reason a rule exists. Kidd's study of knowledge workers points out that the enduring value of such work often lies not in the documents it produces, but in how the work changes the worker (Kidd 1994). When AI generates the artefacts and skips the person, it can flood a company with output while starving it of the learning that output used to leave behind.[1]

[1] Sontag's warning about overproduction is related: a society that produces more than it can attend to dulls its own senses (Sontag 1966).

One natural response to a shrinking present is to keep more of it: record everything, and the past should stay within reach. But abundance of record is not usefulness. Even when the past is stored accurately and in full, its bulk does little to help us act now; it becomes, in a word, *museumified*. As Hoskins observes, possessing the record is very different from being able to find, understand and use it, and it is often the overproduction of the archive itself that hampers access (Hoskins 2024). The recorded past turns at once infinitely searchable and largely inert, kept and catalogued but walked past like exhibits in a museum.[2]

2.3. The Present: Desynchronisation

John Urry argued that a society runs on several kinds of time. Against the shared *clock time* of the industrial era, he set an *instantaneous time*: the always-on, real-time tempo of digital and global flows, in which the collective timetable dissolves and each of us is drawn into a privately managed stream of the moment (Urry 2000). Instantaneous time is desynchronising by its nature: when everything is available at once and on demand, less and less has to be done by everyone together, at the same moment.

With AI, each worker may accelerate inside a private bubble, going faster alone but further from the team (Sarkar 2026). We can already see the next step: agents that attend meetings on our behalf, and summarise those meetings for other agents, humans increasingly drawn to the work of overseeing agentic work – an arrangement that is made more efficient by reducing the "bottlenecks" of shared time and deliberation. As a consequence, shared work routines fragment. Different people and teams come to inhabit different tempos at once, so that a single organisation runs on several clocks that may not agree.

Whose clock the others must keep is a question of power. Sarah Sharma's ethnographies of everyday time, which she gathers under the name *power-chronography*, show that the fashionable story of one universal "speed-up" conceals the fact that no one accelerates on their own, because someone else is always kept waiting, made available, or recalibrated so that the fast can stay fast (Sharma 2014). The frequent-flyer executive keeps to schedule only because an entire architecture of time maintenance, from airport lounges to hotels engineered against jet lag, is arranged around them. The taxi driver waits at the kerb, staying available so that the passenger need never lose a minute. Corporate yoga classes and mid-morning coffee breaks do not free the sedentary worker so much as retune the body for another day at the desk. The slow-food and slow-living movements, which present themselves as refusals of speed, turn out to be temporal arbitrages available mainly to the privileged. Sharma's conclusion is that speeding up and slowing down are both ways of managing people in time, and the point carries over directly to AI: when one worker speeds up on a certain task, time may be redistributed to other parts of the workflow and to other workers.

2.4. The Future: Hyperabundance and the Scarcity of Intention

If the value of the past thins, and the experience of the present fragments, what happens to the future? In many ways, the optimistic vision of accelerated work appears to be a bet on a vastly expanded and enriched conception of the future. AI makes possible futures cheap and plural: ten thousand drafts, designs and directions, all obtainable at once, a hyperabundance of futures. Ought this not to feel like freedom?

Rosa provides an original theory for why abundant futures do not satisfy. Acceleration, he argues, is a secular bid against finitude: *"if we live twice as fast, we can realize twice as many worldly options within a single lifespan; that is, we can live two lives in one"*, so that acceleration becomes *"a functional equivalent to eternity"* (Rosa and Scheuerman 2009). But the promise is ever-receding, because every technique for realising options faster multiplies the options faster still. This is a Jevons paradox of living, in which efficiency at realising futures only whets the appetite for more, so that *"our share of the world, the proportion of realized world options from potentially realizable ones, decreases . . . no matter how much we increase the pace of life"*.

When producing a future becomes cheap and parallel, the scarce resource is no longer production but intention: knowing which of the now-abundant futures is worth having (Sarkar 2024). Intention does not scale the way production does. Confronted with unlimited cheap options, the capacity to want one of them in particular can paralyse. This is one face of the heightened *metacognitive* demand of working with generative AI systems: the burden moves onto knowing what we want, and recognising a good answer when we see it (Tankelevitch et al. 2024). Rosa generalises this diagnosis to society as whole, that acceleration ends in a "hyperaccelerated standstill" (an idea he associates with Virilio and Baudrillard), a frenetic stasis in which, for all the motion, real change is no longer possible (Rosa and Scheuerman 2009).

[2] One could go further: if we rely on AI to retrieve and access the past, the past it returns may not actually have happened. Hoskins notes that such systems can "untether the human past from the present" (Hoskins 2024). In an organisation this need not mean deliberate forgery, but as a consequence of convenience, plausible rationales can be generated for decisions no one holds in organic memory. But that is a secondary concern that we do not press here.

An abundance of futures may paralyse, but it may also devalorise any individual future. In Kundera's *The Unbearable Lightness of Being* (Kundera 1984), what gives a life meaning, per Kundera, is weight: commitment, irreversibility, the road that cannot be retravelled. A future (a document, an analysis, a piece of code) that can be regenerated at no cost carries no such weight: *"the absolute absence of a burden causes man to be lighter than air, to soar into the heights, take leave of the earth and his earthly being, and become only half real, his movements as free as they are insignificant."*

Sontag observed that the camera gave us an over-abundance of *memories*, more captured pasts than we could assimilate (Sontag 1977); generative AI gives the converse, an over-abundance of *futures*, making everyone a kind of point-and-shoot worker. One case, the point-and-shoot programmer, is particularly well exemplified through the practice lately christened "vibe coding" (Sarkar and Drosos 2025).

If AI hands us expected (or expectable) futures for free, then the genuinely surprising futures, out at the edge of the distribution, become especially valuable. Does the practice of pointing the machine and selecting from what it returns develop intentionality and imagination[3] in the same way as the craft practice of photography eventually came to be understood as developing a new "way of seeing"? Photography, Sontag also noted, is at once an art and a utility; the demands of these modes on intention and imagination are entirely different. AI, taken as an engine of culture and work, carries the same duality.

2.5. Power

Virilio's dromology tells us that to understand power, we should follow the speed (Virilio 2006). Three relationships between speed and power are worth pulling apart.

First, speed *is* power. Virilio's image is the naval "fleet in being": not a fleet that gives battle, but *"the permanent presence in the sea of an invisible fleet able to strike no matter where and no matter when"* (Virilio 2006). Its power is entirely potential. By keeping the adversary unsure where the blow might land, it works, in Virilio's words, by *"annihilating the enemy's will to power by creating a global zone of insecurity"* in which the other side can no longer decide with any certainty. This is *"a new idea of violence that no longer comes from direct confrontation and bloodshed"*: the sheer capacity for fast movement dominates without a shot being fired. With AI, a team that can spin up a hundred prototypes on demand, or a firm that could automate a rival's niche overnight, dominates by its credible readiness to act rather than by acting, holding others in a zone of insecurity where no plan is safe from being outpaced.

Second, speed *concentrates* power. As warning times collapse, decisions migrate into fewer and fewer hands. Virilio's example is the Cuban missile crisis: the two superpowers had fifteen minutes' warning time for war, and the Russian rockets on Cuba threatened to cut this down to thirty seconds (Virilio 2006). The fix was to install a direct "hot line" linking the two heads of state. The slow apparatus of war, its staffs, cabinets and legislatures, collapses into the two heads of state themselves, reacting in real time. Virilio then follows the trend to its endpoint: as warning shrinks *"to less than one fatal minute – thus finally abolishing the Head of State's power of reflection and decision in favor of a pure and simple automation of defense systems. The decision for hostilities would then belong only to several strategic computer programs"*. Speed here does not merely centralise power; past a certain velocity it removes the human from the decision altogether. The same logic applies to ordinary organisations: as agents draft, approve and act at machine speed, the human deliberation meant to govern them becomes too slow to accommodate, and oversight calcifies into rubber-stamping. Authority concentrates in the few who set the system's objectives, and in the system that executes them.

Third, speed systems *clash*, and, as an important counterpoint to the first, the fastest side does not always win. In his study of military tempo, Herfried Münkler argues that speeding up and slowing down are opposed powers with different roots: *"the ability to speed things up essentially grows out of economic and especially technological superiority; the ability to slow things down out of the population's capacity for sacrifice and suffering"* (Rosa and Scheuerman 2009). A weaker side can therefore prevail by refusing the fast war and protracting a slow one, whether by waiting out a siege or by fighting a guerrilla war. Münkler's exemplary case is Vietnam, *"the paradigmatic collision of the acceleration and deceleration potentials of two powers"*, won by the decelerating side *"because its opponent's ability to accelerate could not make up for its own population's limited willingness to sacrifice over the long term"*. A fast power also runs on a form of governance that has to keep renewing its own political will, and a long war targets that exact weakness. As Münkler puts it: *"By defensively slowing down the course of military events, the weaker party stood a real chance of resisting far superior attackers and weakening them such that they lost their political will."* A democracy on a four-year electoral clock cannot easily sustain a war of indefinite length, whereas an adversary with no such clock can simply outlast it. Governance and

[3] Imagination, as Benjamin insists, is a contested political capacity rather than a neutral technical one (Benjamin 2024).

war are coupled speed systems, and here the faster one is the more fragile.

For knowledge work this raises an interesting possibility for organisations. Those who cannot match the pace of machine output can decline the speed contest altogether and compete on the slow ground where AI is weakest: judgement, trust, taste, accountability, and socially-embedded relationships. The fast side's flood of output then works in the slow party's favour, and its dependence on speed becomes its Achilles' heel.

2.6. Risks: Accidents and Slowdowns

Virilio observed that every technology of speed is accompanied by a new type of speeding "accident": invent the ship, plane, or automobile and you invent the shipwreck, plane crash, and pileup (Virilio 2006). So what accident does the speed of AI invent? There are of course risks of the loud and spectacular variety, such as the runaway feedback loops of automation too fast for human reaction that we already know from high-frequency trading.

But there are also risks of the less-spectacular variety that are nonetheless tremendously destructive in their cumulative effect: namely, the paradoxical slowdowns that speeding up can produce. An example of this is the traffic jam, where everyone accelerates until the whole system seizes, an outcome no one planned and everyone took part in. In AI-assisted knowledge work, generation is free while judgement stays expensive. This recalls one of Bainbridge's "ironies of automation": the more the machine produces, the more the human is left with the slow, hard residue of watching, checking and verifying it (Bainbridge 1983; Simkute et al. 2025). Drafts, analyses and code arrive faster than human minds can process them, a long line of pull requests appears that no one has time to review, and low-grade output, known by the fashionable epithet "workslop", redistributes the burden of work downstream from its producer (Niederhoffer et al. 2025). The result is knowledge work gridlock.

3. RESEARCH DIRECTIONS

Artificial Intelligence is a chronotechnology, and a research agenda follows from appraising its effects on social time.

One empirical direction is to study how knowledge workers' sense of time actually shifts as they come to work alongside increasingly autonomous systems. Do the past, present and future reconfigure in the ways we have conjectured, and for whom?

Will a glut of possible futures bring on a kind of paralysis as each choice loses its weight? Or might it be felt as the chance to live many lives within one, the condition Camus's "actor" embodies in *The Myth of Sisyphus*, exhausting a plurality of lives in the span of a single one (Camus 1955)?

And what of power: if speed concentrates control, will its gains accrue overwhelmingly to a few players? The felt experience of acceleration is likely not going to be evenly distributed – can we map that unevenness?

We might ask how we can design alternative idioms to speed to conceive of the role of AI in work. Attempts to slow down, from the Slow movement to the "digital detox", deserve study as natural experiments in resistance, with an eye to where they succeed and where they fail. This can build on prior design research on "slow technology" and "temporal design", which asks how artefacts might be built to support reflection and to hold several rhythms at once (Hallnäs and Redström 2001; Pschetz and Bastian 2018). Beside deliberate slowing sit other moves: adaptation, and the cultivation of what Urry called "glacial time", the long, cross-generational pace against which fast time can be measured, and perhaps resisted (Urry 2000).

Finally, the questions Sontag asked of photography can be asked, just as productively, of automation (Sontag 1977). A lot of ink has already been spilled about truth and beauty in machine-made work; about the "interesting" as an aesthetic category; about high culture and low; about creativity (Sarkar 2023), involvement, alienation and detachment; etc. But far less has been said about the intensity of experience; and about the moral position of the automator, who, like the photographer, stands in a relation of capture and appropriation to the futures being automated.

4. CONCLUSION

Jet travel, for those of us fortunate enough to experience it, induces through its instantaneity, through its annihilation of space, a kind of vertigo, in which the body arrives before the mind. AI-assisted work produces the same sensation at the desk.

We have argued that AI can be understood as a chronotechnology; a technology of time, and that the philosophy of speed, acceleration and immediacy is a rich and underused resource for making sense of it. By importing some of the most basic observations from prior thought on social acceleration, we have suggested that AI thins the authority of the past, fragments and desynchronises the present, swells the future into an abundance that makes intention scarce, gathers power along the axis of speed, and invents its own accidents and slowdowns.

We have brought the discourse on artificial intelligence into contact with the philosophies of speed, time, and acceleration. Now we must bring the empirical methods of our field to bear on the study of AI acceleration as an object in its own right.

AI USE DISCLOSURE

The following disclosure statement was generated by the authors using the DAISY protocol (Ahmetoglu et al. 2026): The authors used generative AI tools to improve language and readability in some parts of the manuscript and draft or rewrite text in some parts of the manuscript. All outputs were reviewed and verified by the authors.

REFERENCES

Adams, H. (1918). *The Education of Henry Adams*. Boston: Houghton Mifflin. Chapter "A Law of Acceleration".

Ahmetoglu, Y., M. Constantinides, and A. Cox (2026). Ai disclosure with daisy. In *Proceedings of the 5th Annual Symposium on Human-Computer Interaction for Work*, CHIWORK '26, New York, NY, USA. Association for Computing Machinery.

Bainbridge, L. (1983). Ironies of automation. *Automatica 19*(6), 775–779.

Bauman, Z. (2000). *Liquid Modernity*. Cambridge: Polity Press.

Benjamin, R. (2024). *Imagination: A Manifesto*. New York: W. W. Norton.

Bergson, H. (1910). *Time and Free Will: An Essay on the Immediate Data of Consciousness*. London: Swan Sonnenschein & Co. Translated by F. L. Pogson; New York: Macmillan.

Camus, A. (1955). *The Myth of Sisyphus and Other Essays*. New York: Alfred A. Knopf. Translated by Justin O'Brien; originally published 1942.

Hallnäs, L. and J. Redström (2001). Slow technology – designing for reflection. *Personal and Ubiquitous Computing 5*(3), 201–212.

Han, B.-C. (2015). *The Burnout Society*. Stanford, CA: Stanford University Press. Translated by Erik Butler.

Harvey, D. (1989). *The Condition of Postmodernity: An Enquiry into the Origins of Cultural Change*. Oxford: Blackwell.

Honoré, C. (2004). *In Praise of Slowness: How a Worldwide Movement Is Challenging the Cult of Speed*. San Francisco: HarperSanFrancisco.

Hoskins, A. (2024). Ai and memory. *Memory, Mind & Media 3*, e18.

James, W. (1890). *The Principles of Psychology*. New York: Henry Holt.

Kant, I. (1998). *Critique of Pure Reason*. Cambridge: Cambridge University Press. Translated by Paul Guyer and Allen W. Wood; originally published 1781.

Kidd, A. (1994). The marks are on the knowledge worker. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '94)*, New York, pp. 186–191. ACM.

Kundera, M. (1984). *The Unbearable Lightness of Being*. New York: Harper & Row. Translated by Michael Henry Heim.

Landauer, T. K. (1995). *The Trouble with Computers: Usefulness, Usability, and Productivity*. Cambridge, MA: MIT Press.

Lübbe, H. (2009). The contraction of the present. In H. Rosa and W. E. Scheuerman (Eds.), *High-Speed Society: Social Acceleration, Power, and Modernity*, pp. 159–178. University Park, PA: Pennsylvania State University Press.

McLuhan, M. (1964). *Understanding Media: The Extensions of Man*. New York: McGraw-Hill.

McTaggart, J. E. (1908, 01). I.—the unreality of time. *Mind XVII*(4), 457–474.

Niederhoffer, K., G. R. Kellerman, A. Lee, A. Liebscher, K. Rapuano, and J. T. Hancock (2025, September). Ai-generated "workslop" is destroying productivity. *Harvard Business Review*. Published 22 September 2025.

Poincaré, H. (1913). *The Foundations of Science*. New York: The Science Press. Translated by George Bruce Halsted; includes "The Measure of Time".

Pschetz, L. and M. Bastian (2018). Temporal design: Rethinking time in design. *Design Studies 56*, 169–184.

Rosa, H. (2003). Social acceleration: Ethical and political consequences of a desynchronized high-speed society. *Constellations 10*(1), 3–33.

Rosa, H. (2013). *Social Acceleration: A New Theory of Modernity*. New York: Columbia University Press. Translated by Jonathan Trejo-Mathys.

Rosa, H. and W. E. Scheuerman (Eds.) (2009). *High-Speed Society: Social Acceleration, Power, and Modernity*. University Park, PA: Pennsylvania State University Press.

Sarkar, A. (2023). Exploring perspectives on the impact of artificial intelligence on the creativity of knowledge work: Beyond mechanised plagiarism and stochastic parrots. In *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work (CHIWORK 2023)*, New York, pp. 1–17. ACM.

Sarkar, A. (2024). Intention is all you need. In *Proceedings of the 35th Annual Conference of the Psychology of Programming Interest Group (PPIG 2024)*. arXiv:2410.18851.

Sarkar, A. (2026). Filter Babel: The Challenge of Synthetic Media to Authenticity and Common Ground in AI-Mediated Communication. ACM CHI 2026 Workshop on Restoring Human Authenticity in AI-Mediated Communication (CHI '26 AI-MC).

Sarkar, A. and I. Drosos (2025). Vibe coding: Programming through conversation with artificial intelligence. In *Proceedings of the 36th Annual Conference of the Psychology of Programming Interest Group (PPIG 2025)*. arXiv:2506.23253.

Scheuerman, W. E. (2004). *Liberal Democracy and the Social Acceleration of Time*. Baltimore: Johns Hopkins University Press.

Schmitt, C. (2009). The motorized legislator. In H. Rosa and W. E. Scheuerman (Eds.), *High-Speed Society: Social Acceleration, Power, and Modernity*, pp. 65–76. University Park, PA: Pennsylvania State University Press.

Sharma, S. (2014). *In the Meantime: Temporality and Cultural Politics*. Durham, NC: Duke University Press.

Simkute, A., L. Tankelevitch, V. Kewenig, A. E. Scott, A. Sellen, and S. Rintel (2025). Ironies of Generative AI: Understanding and Mitigating Productivity Loss in Human-AI Interaction. *International Journal of Human–Computer Interaction 41*(5), 2898–2919.

Sontag, S. (1966). *Against Interpretation and Other Essays*. New York: Farrar, Straus and Giroux.

Sontag, S. (1977). *On Photography*. New York: Farrar, Straus and Giroux.

Tankelevitch, L., V. Kewenig, A. Simkute, A. E. Scott, A. Sarkar, A. Sellen, and S. Rintel (2024). The metacognitive demands and opportunities of generative ai. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, New York, pp. 1–24. ACM.

Tomlinson, J. (2007). *The Culture of Speed: The Coming of Immediacy*. London: SAGE.

Ulloa, M., J. L. Butler, S. Haniyur, C. Miller, B. Amos, A. Sarkar, and M.-A. Storey (2026). Product Manager Practices for Delegating Work to Generative AI: "Accountability must not be delegated to non-human actors". In *Proceedings of the IEEE/ACM 48th International Conference on Software Engineering: Software Engineering in Practice*, ICSE-SEIP '26, New York, NY, USA, pp. 10–21. Association for Computing Machinery.

Urry, J. (2000). *Sociology Beyond Societies: Mobilities for the Twenty-First Century*. London: Routledge.

Virilio, P. (2006). *Speed and Politics*. Los Angeles: Semiotext(e). Translated by Mark Polizzotti; originally published 1977.

Wajcman, J. (2015). *Pressed for Time: The Acceleration of Life in Digital Capitalism*. Chicago: University of Chicago Press.

A. THE QUESTIONS OF TIME AND SPEED

Before we ask what AI does to time, or how to reason about the effects of AI acceleration, we need a feel for some of the questions that have preoccupied philosophers of speed. These are broad questions that have been treated by multiple thinkers, and in some sense serve as the "engine" propelling scholarly discourse in this discipline. This section is not a literature review; it is a partial map of the vast territory of thinking about time, speed, and social acceleration that we believe is germane to the accelerating effects of AI.

What is time, and what is it to move through it? This is a fundamental, metaphysical concern. Henri Bergson separated the lived, flowing *durée* of consciousness from the spatialised, clock-measured time of physics, and located our freedom in the first (Bergson 1910). Henri Poincaré, working on the physics of simultaneity, decided that whether two distant events happen "at the same time" is partly a matter of convention: any signal we use to synchronise two clocks takes time to travel, so no experiment reveals an absolute simultaneity, and the astronomer dating a far-off event simply agrees on a rule for the delay (Poincaré 1913). McTaggart went further still: time as we ordinarily picture it, he argued, is contradictory, and therefore unreal. The death of Queen Anne is first future, then present, then further and further past; yet being future, being present and being past cannot all hold at once, and the ordering that does stay fixed, that one event is simply earlier or later than another, seems to leave change out (McTaggart 1908). Two more puzzles sit alongside these: the "specious present", the felt now that itself has some width (James 1890), and Kant's thought that time is not a feature of the world but a form the mind lays over it (Kant 1998).

One high-level conclusion we can draw from all this is that, in order to understand the human experience of time, the idea of an objective, neutral, physical "clock" time is at best irrelevant, and at worst contradictory and unreal. This means that the concepts of speed and acceleration, which in the classical physical sense both rest on a notion of time, ought not to be based on clock time.

For our purposes of analysing the experience of acceleration due to AI, it would seem we ought to adopt a phenomenal approach to speed favoured by theorists such as Tomlinson, more along the lines of the density of experience: *"'Life getting faster' is therefore, most generally, the crowding of incident into our days and the demands this makes upon our resources of energy, time and the attribution of meaning"* (Tomlinson 2007).

How does speed bear on power and politics?
Paul Virilio made speed itself the subject of a discipline he named *dromology*, the study of the race (Virilio 2006). One of Virilio's core claims is that as speed rises, freedom falls. The reason for this is that freedom needs time: time to deliberate, to weigh alternatives, to refuse. Past a certain velocity that time vanishes; decisions must be taken on reflex rather than reflection, and whoever sets the tempo effectively sets the outcome. Speed, for Virilio, is the oldest privilege of the hunter and the warrior, and every advance in it hands power from those who must pause to think to those who can act without pausing. We return to this in Section 2.5.

Acceleration can have similar implications for freedom in institutions. Writing in 1950, Carl Schmitt described the "motorized legislator": under a permanent sense of emergency, from war to inflation to demobilisation, the slow craft of parliamentary lawmaking gives way to fast executive decree, what he called "motorized law", issued faster than legal scholarship can interpret it or a legislature can properly debate (Schmitt 2009). His case in point was the Weimar Republic's slide, through the 1920s and early 1930s, into rule by presidential emergency decree under Article 48 of its constitution, with parliament ceding authority to the executive for no better reason than that the executive was quicker. William Scheuerman builds this into a fuller account of liberal democracy under acceleration (Scheuerman 2004): democratic self-government is deliberately slow, since deliberation, the balancing of interests, and public scrutiny all take time, and that slowness makes its decisions legitimate and revisable. As the economy and technology accelerate, democratic institutions fall out of step with them and begin to lose their grip, and the recurring temptation is to close the gap with something faster: an emergency power, an executive order, or an automated rule.

How do media and technology reshape our relationship to time?
Within media theory, a thread pertinent to our concerns asks how media or technologies change our sense of time. Susan Sontag's essays on photography show how a technology of capture rewrites our relationship to the past, turning lived moments into an ever-growing archive to be browsed (Sontag 1977). Marshall McLuhan argued that electric media collapse distance and remake the senses, so that we take in the world as a simultaneous "mosaic" rather than a linear sequence (McLuhan 1964).

What is social acceleration, and what are its consequences?
John Tomlinson, looking at culture, argues that the old imagery of machine speed has given way to "immediacy": a world of instant delivery, the message answered at once, the film streamed on demand, where the gap between wanting and having seems to close (Tomlinson 2007). Zygmunt Bauman's phrase "liquid modernity", in which things that once seemed solid, from careers to commitments, turn fluid and provisional, also evokes the weightless nature of culture at modern speeds (Bauman 2000). These are broad characterisations of how the speed of modernity has reshaped both culture and individual lived experience. They are not, however, a theory of social acceleration itself. For that we turn to Hartmut Rosa.

Rosa pulls apart three phenomena that are often conflated (Rosa 2013, 2003). First, acceleration in modernity has involved *technological* acceleration: faster movement (the railway replaces the carriage), production (the factory replaces the cottage) and communication (the email replaces the letter). Second, there is the acceleration of *social change*: jobs, institutions and relationships turn over faster. Third, there is the acceleration of the *pace of life*: we feel we have less time and we hurry more, we crowd more incident into our days, even as our devices promise to save it.

Technological acceleration, the acceleration of social change, and the acceleration of the pace of life feed one another in a loop. The irony is that the first was meant to relieve the third and has instead deepened it: the time-saving machine arrives together with a chronic sense of time-famine. Rosa is not the first to observe this irony, but his account provides clarity, as well as an original theory for why the pace of life appears to continually out-accelerate time saving technologies (we will encounter this in Section 2.4).

How might we respond?
In some cases, speed and acceleration are positive. In others, they are detrimental. So what, if anything, should we do about it? Carl Honoré describes a "Slow" movement that wants to win back a humane pace (Honoré 2004). Rosa points instead to *resonance*: not slowness for its own sake, but a responsive, non-instrumental relationship to the world. Tomlinson doubts that any amount of slowing down in private can beat a tempo set for the whole system. Judy Wajcman observes that machines do not set our pace on their own, and how acceleration feels depends on how work, and the sharing-out of

its rewards, is arranged (Wajcman 2015). We return to these answers later. For now it is enough to note that “just slow down” is not the easy answer it appears to be.