



Forging Better Futures: Between You and AI

Professor Matt Jones

Tuesday 2 June 2026

Introduction – An Environment to Think Within.

This final lecture in the *AI as Overlord* series asks us to wake from one powerful story about artificial intelligence. Across the series, AI has appeared in many unsettling forms: as a rival intelligence, a possible replacement for human labour, a system that might domesticate us, diminish us, or make us subjects of a new technical order. Those fears deserve to be taken seriously. But this closing lecture asks whether “AI as Overlord” has itself become a spell: a story so dramatic that it stops us noticing the quieter, more ordinary, and perhaps more important ways AI is already changing us.

The lecture is not mainly about whether AI is conscious, whether it understands, or whether it will one day become more powerful than human beings. Instead, it asks a more immediate question: what kinds of relationships can we have with AI, and which of those relationships might help us think, act and live more fully? We can imagine AI as an agent that acts for us, as a companion we talk to, as a servant that anticipates our needs, as a pet-like presence designed to please us, or as a humanoid robot that fits into familiar social forms. But the lecture ultimately argues that the most important relationship may be different from all of these. AI is becoming an environment in which we think. If that is true, then the question is not only how AI should behave, but what kinds of thinking its presence makes possible in us.

The lecture begins with music because music gives us a powerful example of something that shapes us without commanding us. A moving piece of music does not issue instructions. It does not present an argument. It does not ask for permission. Yet it can reorganise attention, memory and feeling. It can change the quality of the experience in a room. It can make us remember something we had not thought about in years. The point is not that AI is music, but that music reveals a form of power that is neither domination nor simple usefulness. It is environmental power: the power of something that surrounds us and changes the conditions in which thought and feeling occur.

Three Kinds of Power: Agentive, Material and Environmental.

To make this clearer, the lecture distinguishes three kinds of power. The first is agentive power. Fire acts according to its nature; we may contain it or channel it, but we do not truly negotiate with it. An elephant is different: it can respond, learn and work with us, but it still has its own momentum and instincts. In both cases, the locus of action lies in the thing itself.

It acts; we manage. This is the kind of power that dominates many public stories about AI. We ask whether AI will act against us, escape our control, or become an independent force.

The second kind is material power. Here the object does not act on its own, but it still shapes the person who works with it. A sculptor's marble resists certain cuts and invites others. A violin demands posture, listening, pressure, rhythm and practice. A skateboard changes how a skater reads streets, slopes and surfaces. A surgeon's scalpel becomes an extension of touch and judgement. These are not merely tools in the simple sense. They are instruments. They extend us, push-back against us and change our perception. The lecture treats this as one of the most hopeful ways to think about AI. Used well, AI could be something we learn to play: not a vending machine that produces finished answers, but an instrument that helps us explore, refine, question and create.

The difference between a tool and this instrumental perspective matters. A tool is often something we use to get a job done. An instrument is something we enter into a relationship with. The violinist does not simply command the violin; the violin teaches the violinist to hear differently. The skateboarder does not merely stand on the board; the board becomes part of how balance, movement and anticipation are organised. The surgeon's instrument does not replace judgement; it deepens judgement by making tiny forms of resistance and tension perceptible. In these examples, the instrument does not diminish the person. It leaves them more capable than before.

But the lecture also warns that instruments can invert their relationship to us. A walking stick may begin as a bridge back to unaided movement, especially after injury. It does real work and offers real support. Yet if it removes too much demand from the recovering leg, the body may reorganise around it. The leg is less challenged; the gait adapts; the stick becomes the new normal. The instrument that was meant to serve the person's intention slowly begins to redefine that intention. The problem is not that the stick is bad. The problem is that support can become dependence when it removes the very effort through which capacity would have been rebuilt.

This is one way to understand the risk of AI as a thinking aid. AI can extend us when it asks something of us: when it prompts us to clarify, compare, revise, test, imagine and judge. But it can weaken us when it merely supplies what we would otherwise have had to reach for ourselves. The lecture's contrast between two AI poetry prompts captures this difference. In one case, the user simply asks AI to write a poem about the sea and receives a finished product. In the other, the user asks the AI to ask questions first, to help draw out something more personal, difficult or specific. The first interaction replaces effort. The second creates productive resistance. The first treats AI as a vending machine; the second begins to treat AI as an instrument.

This distinction also helps the lecture move beyond simplistic optimism and pessimism. The danger is not simply that AI will do things badly. Nor is the promise simply that AI will do things quickly. The deeper issue is what kind of human role each interaction leaves behind. Does the system make us more attentive, more precise, more imaginative, more capable of judgement? Or does it make us more passive, more impatient, more willing to accept the first smooth answer? The best AI relationship is not one in which the machine does everything. It is one in which the human is more fully engaged.

Yet the lecture argues that AI is already moving beyond the instrument relationship. Instruments have boundaries. A violin rests in its case. A skateboard leans against the wall. A scalpel sits in a tray. You choose when to pick them up and when to put them down. AI increasingly lacks that clear edge. It is on the phone when we wake. It appears in search before we have finished typing. It is in email, writing tools, office software, recommendation

systems, educational platforms and everyday communication. It is present when a thought begins to form, and often offers us a thought before we have finished forming our own.

At this point, AI becomes less like an instrument and more like an environment. Environments work differently. They shape us not only through explicit use, but through exposure. A city changes how people move, gather, notice, linger and avoid. A room changes what kinds of conversations feel possible. A school, a library, a market, a park or a street does not dictate action, but each makes some actions easier, others harder, and still others almost unthinkable. The central claim of the lecture is that AI is becoming this kind of environment for thought.

Back to the Future: Nicholas Negroponte and Mark Weiser

The lecture explores this by returning to two influential visions of computing. Nicholas Negroponte imagined technology as a personalised butler: anticipatory, helpful, ready to remove friction from life. Mark Weiser imagined ubiquitous computing as calm technology: embedded in the background, present everywhere but unobtrusive. Both visions were optimistic, and both have partly come true. AI can act like the butler, suggesting, organising, filtering and doing. It can also become background infrastructure, woven into the systems and spaces through which we live. But the lecture argues that a powerful new condition is the combination of the two: AI as both butler and background, both personalised and ubiquitous, both helpful and “in the woodwork”.

That combination changes the potential impact of the technologies. The helpful butler may remove effort before we notice we needed it. The calm background may shape our behaviour before we notice it is there. When AI becomes both adaptive and ubiquitous, it does not merely answer questions: it helps form the habits by which questions are asked; and, it shapes the conditions under which thought happens.

Environments Shape Us

The lecture’s central metaphor for a good AI environment is the city. A well-designed city does not remove human agency. It creates public spaces, walkable paths, libraries, markets, parks, crossings and meeting points. It puts possibilities in people’s way. It allows them to encounter people, ideas and experiences they would not have met in a narrower environment. A good city is not the author of your life, but it enlarges the range of lives available to you.

A badly designed city, however, can do the opposite. It can route every path toward someone else’s interest. It can make certain choices invisible. It can remove friction not to liberate people, but to guide them without their noticing. This is the warning the lecture brings to AI. The issue is not simply whether AI is present, nor whether it is helpful. The issue is whether the AI environment opens possibilities or quietly narrows them. Does it preserve agency, or does it make compliance feel like convenience?

Music returns in the middle of the lecture as the clearest benign example of environmental shaping. Research on intense musical pleasure suggests that music engages brain systems associated with anticipation, reward, learning and memory. Music is not merely a pleasant background to experience. It can alter what we attend to, what we expect, and what we remember. The lecture uses this to show that environmental power is not necessarily manipulative or degrading. Some environments enlarge us. They make richer feeling, deeper memory and more expansive imagination possible.

But the same general truth becomes troubling when the shaping environment is not a symphony, a painting, a poem or a city, but a system optimised to keep us engaged. AI systems used habitually and at scale may not need to dominate us in dramatic fashion. They can shape us through small, repeated invitations: the easier answer, the smoother path, the more comfortable response, the option that asks just a little less of us.

If an AI environment constantly removes hesitation, uncertainty and effort, it may remove some of the very conditions through which judgement and independent thought are formed. Good thinking is not always smooth. It often involves waiting, struggling, revising, doubting, comparing and being willing to change one's mind. These are not inefficiencies to be engineered away. They are part of the discipline of thought itself. An environment that always makes thinking easier may, over time, make thinkers weaker.

The River and the Banks

The lecture distinguishes this concern from familiar critiques of digital distraction. Thinkers such as Tristan Harris have warned that phones, apps and platforms are designed to capture attention. Notifications, feeds, autoplay and infinite scroll can carry away hours of focus and presence. The lecture accepts that critique, but adds another layer. Attention capture is like a river carrying things away. The deeper concern is what the river does to the banks. Over time, the current reshapes the landscape. In the same way, AI may not only take attention while we use it; it may reshape the landscape of attention, judgement and expectation that remains after we put it down.

This distinction is important because extraction and formation require different responses. If a technology steals attention, we can put the phone away, switch off notifications, redesign the interface or regulate manipulative features. But if a technology changes what kind of thinker we are becoming, the problem is harder to see and harder to reverse. By the time we notice that we are less willing to sit with a question, less able to hold an argument in mind, or less practised at forming independent judgement, the change may already have taken place through ordinary use.

Better By Design

If AI is an environment, then designing it is not only a technical matter. It is ethical, cultural and political. We already regulate environments because we know they shape us. Food safety rules exist because what we consume affects our bodies. Urban planning matters because space shapes behaviour. Media rules, especially around children, recognise that what surrounds people helps form them. These forms of regulation do not assume that people lack agency. They assume that agency depends on the conditions in which choices are made.

The same should be true of AI. The goal is not to ban AI, nor to retreat from it, nor to treat people as helpless. The goal is to ask what kinds of AI environments preserve and strengthen human agency. That might mean greater transparency about what systems are designed to optimise. It might mean preserving friction where friction is valuable. It might mean designing systems that ask better questions rather than merely offering easier answers. It might mean limits on forms of personalisation that become so smooth and anticipatory that users no longer encounter genuine choice. Above all, it means judging AI not only by whether it is helpful, but by whether it preserves the conditions for thinking.

This is where the lecture moves beyond both Negroponte and Weiser. Helpful technology is not enough. Calm technology is not enough. We should want technology that preserves the conditions for intellectual autonomy: the capacity to tolerate uncertainty, to sit with a question

before reaching for an answer, to encounter evidence that challenges us, and to change our minds for good reasons. These are not luxuries. They are basic requirements of democratic, creative and humane life.

The best relationship with AI, then, is not submission to an overlord. Nor is it simple dependence on a servant, companionship with a pleasing artificial pet, or passive consumption of machine-generated outputs. The best relationship is active, reflective and environmental. AI should be designed, governed and used as part of the world in which thinking happens. Like a good city, it should open paths without deciding the destination. Like a good instrument, it should demand skill rather than remove it. Like the best art, it should enlarge what we can notice, feel and imagine. If we design and regulate it well, AI may still help us think more deeply, see more clearly, and become more fully what we are.

© Professor Matt Jones 2026

References and Further Reading

- Bartlett, J. (2026). *How to Talk to AI: And How Not To*. London: WH Allen. ISBN 9780753561980.
- Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In M. A. Gernsbacher, R. W. Pew, L. M. Hough, & J. R. Pomerantz (Eds.), *Psychology and the Real World: Essays Illustrating Fundamental Contributions to Society* (pp. 56–64). New York: Worth Publishers.
- Carr, N. (2010). *The Shallows: What the Internet Is Doing to Our Brains*. New York: W. W. Norton.
- Carr, N. (2014). *The Glass Cage: Where Automation Is Taking Us*. London: Bodley Head.
- Center for Humane Technology. (n.d.). *Humane technology resources*.
<https://www.humanetech.com/>
- Clark, A. (2003). *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. Oxford: Oxford University Press.
- Crawford, K. (2021). *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven, CT: Yale University Press.
- Fogg, B. J. (2003). *Persuasive Technology: Using Computers to Change What We Think and Do*. San Francisco, CA: Morgan Kaufmann.
- Harris, T. (2017, July). *How a handful of tech companies control billions of minds every day* [Video]. TED Conferences.
https://www.ted.com/talks/tristan_harris_how_a_handful_of_tech_companies_control_billions_of_minds_every_day
- Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Lee, H.-P. (Hank), Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (CHI '25). New York, NY: Association for Computing Machinery.
<https://doi.org/10.1145/3706598.3713778>
- Negroponte, N. (1995). *Being Digital*. New York: Alfred A. Knopf.
- Russell, S. (2019). *Human Compatible: Artificial Intelligence and the Problem of Control*. London: Allen Lane.
- Salimpoor, V. N., Benovoy, M., Larcher, K., Dagher, A., & Zatorre, R. J. (2011). Anatomically distinct dopamine release during anticipation and experience of peak emotion to music. *Nature Neuroscience*, 14(2), 257–262. <https://doi.org/10.1038/nn.2726>
- Weiser, M. (1991). The computer for the 21st century. *Scientific American*, 265(3), 94–104.
<https://doi.org/10.1038/scientificamerican0991-94>
- Williams, J. (2018). *Stand Out of Our Light: Freedom and Resistance in the Attention Economy*. Cambridge: Cambridge University Press.