

# What Does ‘Human-Centred AI’ Mean?

Olivia Guest<sup>1,2</sup>

<sup>1</sup>Department of Cognitive Science and Artificial Intelligence,  
Radboud University, The Netherlands

<sup>2</sup>Donders Institute for Brain, Cognition, and Behaviour,  
Radboud University, The Netherlands

**Abstract:** While it seems sensible that human-centred artificial intelligence (AI) means centring “human behaviour and experience,” it cannot be any other way. AI, I argue, is usefully seen as a relationship between technology and humans where it appears that artifacts can perform, to a greater or lesser extent, human cognitive labour. This is evinced using examples that juxtapose technology with cognition, *inter alia*: abacus versus mental arithmetic; alarm clock versus knocker-upper; camera versus vision; and sweatshop versus tailor. Using novel definitions and analyses, sociotechnical relationships can be analysed into varying types of: displacement (harmful), enhancement (beneficial), and/or replacement (neutral) of human cognitive labour. Ultimately, all AI implicates human cognition; no matter what. Obfuscation of cognition in the AI context — from clocks to artificial neural networks — results in distortion, in slowing critical engagement, perverting cognitive science, and indeed in limiting our ability to truly centre humans and humanity in the engineering of AI systems. To even begin to de-fetishise AI, we must look the human-in-the-loop in the eyes.

**Keywords:** artificial intelligence; cognitive science; sociotechnical relationship; cognitive labour; artificial neural network; technology; cognition; human-centred AI

## 1 Introduction

[M]odels have to be used before they will give up their secrets. In this sense, they have the quality of a technology — the power of the model only becomes apparent in the context of its use. (Morrison & Morgan, 1999, p. 12)

We are experiencing a point in history that considers itself separate, unique, a zenith of human civilisation, presented as disconnected from the past through a series of technological sea changes. In extreme cases, we are detached from even the last few months: the newest version of a piece of artificial intelligence (AI) software is the definitive one, everything else is irrelevant. And so findings in AI are held to expire. Our AI students describe their own knowledge and skills — even their whole degree itself — as having upcoming sell-by dates. Research in AI is thus often framed as unmoored from historical, ethical, social, and legal precedents.

The speed of the “ultra rapid computing machine,” as Wiener (1948, 1950) describes computers, becomes a metaphor for the speed of change and of technological progress. This notion of condensed time operates as a further discursive regularity in two ways. First, it forms an overall temporal backdrop against which various cybernetic dramas are played out. Second, condensed time becomes a measure of the performance of humans and machines.

Hughes (1985) offers a telling example of the first manifestation of condensed time:

In the scant two hundred years since the early Industrial Revolution, Western man has totally rescaled and changed the face and fabric of his environment. Indeed, the changes have proceeded at such an accelerated pace that we might use the word “old” or “outmoded” to refer to last month’s computer model. (p. 205)

This sense of condensing time—of speed, of rapid change—imbues virtually all the popular literature examined. It becomes a contextual operator, stated but not questioned. (Hamilton, 1998, p. 193)

What Sheryl N. Hamilton describes above is more than seven decades old. But if this has been going for decades, two-thirds of a century or longer still, does this not undermine its very premise? How can we both be at a unique point in history, and trapped in an endless cycle that began with the paralleling of humans and clockwork?

The process of uncondensing time indubitably requires slowing down (Stengers, 2018) and radically centring the human and decentring the “ultra rapid computing machine.” (Hamilton, 1998) To do this, some propose to (conceptually) reengineer (Guest, 2024; Nado, 2021) or reimagine these machines and our relationships to them *inter alia* through a subfield known as human-centred AI. For a pertinent example, Andy J. Wills (2025), in the call for this special issue, describes human-centred artificial intelligence (HCAI) as placing “human behaviour and experience at the heart of [AI] research.” An example of this, he says, could be cases wherein it is “claimed that artificial neural networks (ANNs) now perform at human levels in a variety of tasks”. He goes on to ask:

Can ANNs effectively and safely be used to support the work of highly trained professionals—for example, radiologists, therapists, legal advisors, or researchers? Can we effectively adapt the skills and techniques of behavioural research, previously applied to humans and other animals, to better understand the ‘psychology’ of complex black-box ANNs?

Relatedly, Yvonne Rogers (2022) states the goal of HCAI as “designing AI systems that enhance human capacities and improve their experiences rather than replacing them through automation[, along with the aspiration of] reimagining human-machine interaction in all its guises[, while negotiating] the creeping creepiness of AI” (p. 255). Industry players, such as IBM, make a different but not entirely incompatible conceptualisation:

Human-Centered AI (HCAI) is an emerging discipline intent on creating AI systems that amplify and augment rather than displace human abilities. HCAI seeks to preserve human control in a way that ensures artificial intelligence meets our needs while also operating transparently, delivering equitable outcomes, and respecting privacy. (Geyer et al., 2022, n.p.)

Regulatory bodies, like the European Union, also have a perspective:

The human-centric approach to AI strives to ensure that human values are central to the way in which AI systems are developed, deployed, used and monitored, by ensuring respect for fundamental rights, [...] which are united by reference to a common foundation rooted in respect for human dignity, in which the human being enjoys a unique and inalienable moral status. (Tambiana, 2019, p. 3)

Overall we see that HCAI — as a specialisation for students (e.g. “Human-Centered AI track, MSc Human-Technology Interaction, Eindhoven University of Technology”, 2025; “Human-Centred AI theme, MSc Data Science and Artificial Intelligence Technology, TU Delft”, 2025), as a subfield for scholars (e.g. “HCAIM Consortium”, 2025; Holzinger, 2025; “Human-centered Artificial Intelligence, Utrecht University”, 2025; Rogers, 2022; Stanford Institute for Human-Centered AI, 2019; Wills, 2025) or technology industry workers (e.g. Geyer et al., 2022) alike, and as a concern and mission for regulators (e.g. Pirozzoli, 2024; Tambiana, 2019) — interweaves four related themes:

1. supporting or enhancing human skills, both without displacement of said skills and without violation of fundamental human rights;
2. imbuing systems with various so-called human-aligned values, including those of explainability and transparency;
3. focussing on human behaviour as a benchmark, i.e. the idea of human-like or -level performance;
4. implicating behavioural, or otherwise psychological, methods in the study of these systems as if on equal epistemic footing with humans.

In this paper, I turn this frame on its head (cf. Andersen et al., 2023; Bannon, 2011; Bishop, 2021; Nick Dyer-Witheford, 2019; Rogers, 2022; Ryan, 2024; Schmager et al., 2025; Shneiderman, 2022; Steinhoff, 2021; Walton & Nayak, 2021). Using a unique (re)definition of AI that releases us from a correlationist<sup>1</sup> grip, we then examine three triplets as case studies of techno-social relationships between cognition and artifacts. To presage the coming analyses, herein AI is any techno-social relationship that outsources to machines or algorithms some part, however small, of human cognitive labour (cf. van Rooij et al., 2024). I demonstrate that AI is human-centric, *not* because it behaves like or is designed to be like humans, but because it requires a ghost in the machine, often literally an obfuscated human-in-the-loop to properly function (also see Guest & Martin, 2025) because AI *is* humans albeit in fetishised (Braune, 2020; Morris, 2017; Mota & Cosentino Filho, 2024), obfuscated forms (e.g. Erscoi et al., 2023). That is, AI “is in reality produced by relations among people [even though it] appears before us in a fantastic form as relations among things” (Pfaffenberger, 1988, p. 250). AI’s “technological veil” hides human cognitive labour (Mota & Cosentino Filho, 2024). I bring this anthropological, sociotechnical and broadly computational cognitive scientific angle to understanding AI, that is unlike “the skills and techniques of behavioural research, previously applied to humans and other animals[, which claim to help us] better understand the ‘psychology’ of complex black-box ANNs” (Wills, 2025, n.p. cf. Raley and Rhee, 2023). In fact behavioural probing of such systems using such experimental techniques assumes their psychological standing to be equal to, or comparable to, biological organisms: it begs the question (see for relevant analyses: Forbes & Guest, 2025; Guest & Martin, 2023, 2024; Raji et al., 2022).

The perspective I bring here answers questions upstream to analysing the behavioural outputs of such systems, instead focussing on *in principle* analyses, freeing us from a correlationist account which delivers flawed reasoning and uninterpretable results (Guest & Martin, 2023, 2024; Guest, Scharfenberg, & van Rooij, 2025). The contradiction between artifice, artificiality, machines, and the machinic and intelligence, cognition, and humanity is problematised and dissected by the method herein. These two can appear both as opposites and as identical; at odds as analytical constructs and inexorably intertwined. For example, “[f]or a long time, the human was something else altogether; it is not so long ago that it became a machine—a calculating one no less.” (Mauss, 1923, pp. 176–177; translated by LePage-Richer, 2024, p. 20) Some go further along this route, remarking that “the history of the sciences is now reaching a point, in all its branches, where every scientific theory can be taken as a machine” (Guattari et al., 1984, p. 112); and arguing “that scientific models have certain features which enable us to treat them as a technology.” (Morrison & Morgan, 1999, p. 35) And others — coinciding with the regulators’ worries — notice that “[m]achinery does not just act as a superior competitor to the worker, always on the point of making [them] superfluous. It is a power inimical to [them], and capital proclaims this fact loudly and deliberately, as well as making use of it.” (Marx Memorial Library, 2024, n.p.) Which is it? Are humans machines, or machinic in some important way, so they can be recreated in machines? Are machines somehow human-like, created by us in our image? Let me tell you.

## 2 Radically Redefining AI

[I]t is impossible to create an absolutely reliable automatic system, and sooner or later people face the necessity to act after equipment fails. [...] If the cosmonaut loses such skills because of [their] passive role [due to being typically limited to monitoring and observation only], the probability of [their] choosing and carrying out the right procedure in an emergency would be small. This contradiction is inherent in automatic control systems. (Ponomareva, 1998, n.p.)

As practitioners of science, we are duty bound to think about what ‘artificial’ and ‘intelligence’ mean. What does it mean to propose an artificial version of a human capacity? When we talk about redefining something, we must remember that “definitions have no inherent truth. They are agreed conventions tested by their internal coherence as well as by their relationship to common sense, common practice and history.” (Yuval-Davis 2024, p. 790; also see Elgin

<sup>1</sup>The idea that it is correlations with a given target, like human data derived benchmarks (Pasquinelli 2017; Schrimpf et al. 2020; cf. van Rooij et al. 2024), that validate the safety (El-Mhamdi et al., 2022), appropriateness, or other desirable properties of a model or system. In general, correlationism may be unproblematic, but in a setting where computation is involved correlations cannot function as a useful guide, serving more as red herring than anything else (Guest & Martin, 2023, 2024, 2025; Guest, Scharfenberg, & van Rooij, 2025).

Table 1. The two steps required for the proposed redefinition of AI. At the top is step 1, where we decide whether a relationship exists between a technology and human cognition. This relationship, represented by the blue-green column between Machine and Human on row 1a, is AI. In 1b are terminological examples, both non-diagnostic on their own and incomplete as a list, that can aid in the diagnosis of a sociotechnical relationship as one of AI. The three columns below in step 2 represent three, not mutually exclusive, types of sociotechnical relationship between humans and artifacts. At this step, we sketch out if AI replaces, enhances, or displaces cognition (row 2a) — with relevant properties and their typical values, non-exhaustively specified, listed on rows 2b–h.

1) Discern relationship:

Does the artefact relate to cognitive labour?

a) entities involved	MACHINE	AI	HUMAN
b) potential terms	algorithm, artificial, automation, benchmark, discriminative, computer, engineering, functional, generative, mechanism, system, technology, tool		ability, behaviour, capacity, cognition, intelligence, labour, learn, organism, professional, psychology, reason, skill, task, thought, train

2) Characterise relationship:

How does the artefact relate to cognitive labour?

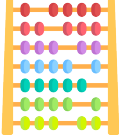





a) label	REPLACEMENT	ENHANCEMENT	DISPLACEMENT
b) valence	neutral	beneficial	harmful
c) effect on cognition	unaffected	reskilling	deskilling
d) labour obfuscation	minimal	possible	maximal
e) human equivalence	worse or same	different	no
f) human-in-the-loop	rare	possible	common
g) human input	transparent	transparent	opaque
h) desired output	specified	formal	unspecified

2017; Nado 2021) Some of the existing definitions of AI when probed reveal weaknesses in the aforementioned dimensions — especially undesirable when speaking of formal systems, which all AI systems necessarily are. This is to say, “weakness from the viewpoint of formal symmetry of doctrine [is] strength in the service of rising capitalism.” (Novack, 1968, p. 25) So these formal and common-sense problems with pre-existing AI conceptions form a core of strength when incentives are, in opposition to academia, ones of profit-making for the technology industry (see van Rooij et al., 2024, for an exposition of possible meanings). AI has had so many related but nonetheless different meanings over the decades that perhaps it is a fool’s errand to even try and pin it down (Boden, 2006; Dreyfus, 1965; Guest & Martin, 2024; Haigh, 2023; Lighthill et al., 1973; McCorduck, 2004; Smith & Smith, 2024). And similar issues with terminology appear with related terms like ‘brain inspired’ and ‘neurally plausible’ to name but a few terms which function as weasel words (Guest & Martin, 2024) — as well as when ‘computational’ is used metaphorically or naively by cognitive neuroscientists (Guest & Martin, 2025; Guest, Scharfenberg, & van Rooij, 2025).

All AI implicates human cognition, as a user, as a human-in-the-loop, and as an inspiration. What I propose is that the important aspect is understanding how it does this; understanding the sociotechnical *relationship*, as opposed to understanding (e.g. probing, benchmarking) the technology in the abstract. The shifting of the scientific emphasis to the relationship — what does the model do as a function of interacting with humans? — from what the model is in and of itself is by no means completely alien to cognitive scientists (Bainbridge, 1983). For example, cognitive computational modellers are familiar, consciously



Table 2. The first column represents the AI relationship (defined in Table 1) between using an abacus and non-abacus assisted mental arithmetic. The second column between using an electronic or mechanical calculator and an unassisted adult who knows basic mathematical operations already. The third column between using a digital computer and a human computer. The photo in the third column depicts NASA human computers Dorothy Vaughan, Lessie Hunter, and Vivian Adair (Shetterly, 2016a, 2016b).

artefact versus cognitive labour	ABACUS VERSUS MENTAL ARITHMETIC	CALCULATOR VERSUS HUMAN	DIGITAL COMPUTER VERSUS HUMAN COMPUTER
			
			
a) label	enhancement	replacement	displacement
b) valence	beneficial	neutral	harmful
c) effect on cognition	reskilling	unaffected	deskilling
d) labour obfuscation	none	moderate	high
e) human equivalence	same	better	quantity, not quality
f) human-in-the-loop	abacist	none	programmer
g) human input	full-blown cognition	function, numbers	code, input to code
h) desired output	result of calculation	result of calculation	result of calculation

or otherwise, with what in the philosophy of science is known as the pragmatic view on scientific theories and models, i.e. models are characterisable by their properties derived from their use (Guest & Martin, 2021; Morgan & Morrison, 1999; Winther, 2021).

Without further ado and notwithstanding terminological disarray, Table 1 embodies a single tripartite definition that acts as an analytical tool for discerning the AI's properties, e.g. labour obfuscation, which is the extent of the hidden human-in-the-loop (Crawford, 2021; Guest & Martin, 2025). The radical redefinition I propose comprises two parts. First, full-blown deflation — any sociotechnical relationship could be AI if it links an artefact to human cognition such that the artefact can be seen as performing some aspect of a cognitive capacity (cf. Egan, 2025; Guest & Martin, 2025). For this step, all we need to ask ourselves is: does the technology exist in a relationship with human cognition? If we affirm this, we can move to the next step. Second, we reinflate AI into three types of relationship, that of enhancement, of replacement, and of displacement (see Table 1) of human cognitive labour.

As mentioned, AI is any relationship between technology, tools, models, machines and humans where it appears as if cognitive labour is offloaded onto such artifacts. Anthropological notions of tools versus technologies can help to unpack what is going on with respect

to such artifacts:

Tools are creations on a localised, small-scale, the products of either individuals or small groups on specific occasions. As such, they do not give rise to systems of control and coercion. Technology, on the other hand, is the product of large-scale interlocking systems of extraction, production, distribution and consumption, and such systems gain their own momentum and dynamic. (Moore, 1997, p. 5)

On this, we can readily see an abacus can be both a tool and exist in a relationship to human cognition. A familiar case of this type of offloading is sending off your desired arithmetic function and the arguments, e.g. adding two numbers, to a digital calculator. In this example, you do not add the numbers yourself. You do three things: know which numbers you need, know which function you need, and know how to use the machine to offload what would otherwise be your cognitive labour. This is step one: yes, it is AI. Had the machine not existed, you would have had to perform the arithmetic yourself, and arithmetic is cognitive labour, therefore, yes. The AI relationship could be any of the three in Table 1, i.e. more information is needed to understand what is going on in a specific use-case of a calculator to perform addition. Let us pick as the human in this relationship, a familiar one to us all, a child who knows the symbols for numbers and addition, and how to use the calculator, but does not yet know mental or otherwise addition. The effect that offloading addition to the calculator has in this case is why we ban their use by children who have not yet learned basic numeracy, i.e. in this case using a calculator embodies a way to avoid learning how to add numbers from the rote learning of the addition of pairs of numbers under ten to the rules for larger numbers. As a society then, we decide it is undesirable and mostly the 3rd column of step 2 in Table 1: *a*) displacement. It is *b*) harmful to the development of *c*) the child's numeracy skills and solely knowing how to add using a calculator is not knowing how to add because problems will appear when, e.g. the numbers to be added are beyond the maximum number representable by the calculator, in such a case *d*) obfuscation of somebody else's cognitive labour is likely to happen as the calculator on its own without somebody who knows addition is not *e*) equivalent to a human.<sup>2</sup>

For a calculator generally, typically the human-in-the-loop, involvement of human oversight or data after the user input, *f*) is either absent or minimal (e.g. to change the batteries) — and relatedly, the input *g*) is easily identifiable as such, i.e. what the user must have done to obtain *h*) the desired output, in this case the result of the arithmetic operation requested, which is formally well-specified and verifiable. Importantly, a different relationship to humans as calculators typically enjoy is not displacement AI, but likely enhancement AI: with calculators we lose none<sup>3</sup> of our extant mathematical skills as adults and enjoy a shorter time complexity on numerical operations (see column 2, Table 2). Not all technology has such a trajectory, as we shall see with specific cases of this in-depth in the next section.

### 3 Artefacts versus Cognitive Labour

There can be no doubt that the idea of “intelligent”, “thinking” machines has captured the imagination of many people all over the world. (Saparina, 1966, p. 295)

AI is unlike a tool like a saw used to cut wood where the person cutting also puts in labour, often more than the creator of the saw, to cut wood, in the sense that a saw is involved in the predominantly overt physical and cognitive labour of woodworking (Guest, Suarez, et al., 2025). Not so for typical cases of contemporary AI, like chatbots — in contrast: You did not contribute other than as input, e.g. the so-called prompt, ultimately harming the chance to learn anything substantial (e.g. Bastani et al., 2024; Guest, Suarez, et al., 2025). It is never clear if the chatbot's results really match those desired, e.g. so-called hallucinations — “a misleading (and anthropomorphizing) term[, which has become mainstream, and thus] a major win for AI hype” (Helfrich 2024, p. 700; Bishop 2021; M. T. Hicks et al. 2024). Due to their opaque nature, AI relationships hide labour that might be invoked not only like the saw to extract raw materials or design systems, but also in real-time as you use the artefact, e.g. you forget that sweatshop workers may be in real-time or the recent past refining the

<sup>2</sup>In the case of adding numbers outside the maximum, a person who does know addition can use the calculator to perform parts of the sum or resort to paper or any other combination of the above.

<sup>3</sup>Perhaps skills like long division atrophy, but we tolerate this.

output of the chatbot (Crawford, 2021; Perrigo, 2023; Placani, 2024). So much for a stark contrast between a tool-human sociotechnical relationship without AI and a relationship that involves modern chatbots, what about technologies that are more complex than a saw?

Below, three triplets are presented for exploring and exemplifying my definition of AI (from Table 1):

**3.1 Abacus, Calculator, Computer:** all three artifacts (in Table 2), which are prototypical computational aids or devices, diverge greatly on their need for direct human involvement, but nonetheless can share the same desired output (the result of a given calculation).

**3.2 Alarm Clock, Camera, Garment Factory:** these three examples (in Table 3) demonstrate how divergent and different the cognitive labour and capacities (knocker-upper, the human alarm clock; human vision; seamstress/tailor) and their related artefacts are.

**3.3 LLM, Image Generator, Chatbot:** these three (in Table 4) implicate what is typical contemporary AI, very deep artificial neural networks trained on extremely large datasets (a correlationist programme previously dubbed *modern connectionism* in Guest & Martin, 2024) with the cognitive labour they claim to capture, of which all relationships are characterised as displacement.

### 3.1 Abacus, Calculator, Computer







Analysing these ancestral forms of computers — which function as aids, in the case of the abacus, and as perhaps something more independent in the cases of the calculator and Turing-complete electromechanical and electrical digital computers — as AI per Table 1 brings to light aspects that are central to understanding cognitive labour. For the abacus versus mental arithmetic AI relationship (see column 1, Table 2), we — surprisingly or not depending on our familiarity with abacus use — see a marked benefit to mental arithmetic and an unambiguous development of a new skill (Lima-Silva et al., 2021; Lu et al., 2023; Wang, 2020; Xie et al., 2024). Importantly, the examples in Table 2 have been chosen *inter alia* because we can keep row *h* constant: all three have the same desired output in these cases. An interesting highlight is row *g*) the abacus is completely unable to perform any arithmetic operation without the abacist. If you forget how to manipulate the beads, the abacus is merely decorative. This is very different to the electronic calculator (see column 2, Table 2) where one need not be familiar with the operation of the calculator at all: you only need to recall the symbols that represent numbers and functions. Not knowing how to use a calculator is nigh on impossible in the modern world, even without having ever used one, since the prerequisite knowledge is reading and button affordances; no specific training is needed. To forget how to use a calculator is to suffer significant cognitive impairment beyond being rusty moving beads around; one would need to lose the ability to read numbers and press buttons — highly unusual. And the opposite, using a calculator can harm children's ability to accomplish basic numeracy, and so we proscribe its use in primary schools for this reason. In the general case, once mental arithmetic is mastered the AI relationship between calculator and human is overall neutral without offering any new skills but also providing predictably speedy and verifiably correct arithmetic. The thermostat is another great example of no deskilling on an individual level — before that we could only sense using our bodies if things were too hot or too cold, and we still can.

In contrast to these positive and neutral AI relations is the original human-computer relationship (see column 3, Table 2). A human computer was a person — often a woman in the previous century, but less gendered prior to that — who performed calculations, worked with computing machinery, wrote programming software (Grier, 2013; Shetterly, 2016b). For example, during Britain's war effort against the Nazis:

Arriving members of the [Women's Royal Naval Service] were given two weeks training in binary math, the teleprinter alphabet, sight-reading punched paper tapes, and the structure and workings of the Tunny and Colossus machines. [...]

Machine work—and the theory and skills it required—was an integral component both intellectually and functionally of [the Second World War's] codebreaking process. It was not, as many assumed due to its feminized nature, deskilled drudge work. (M. Hicks, 2017, pp. 39–40)

Table 3. Three examples (columns) of technosocial systems (e.g. user and alarm clock) paired to what came before (knocker-upper; column 1) or to a classical cognitive capacity (vision; column 2) or to a non-sweatshop version of similar skills and labour (seamstress/tailor; column 3) to demonstrate that their important properties (rows) can be teased apart and understood if we center humans in our analyses (recall Table 1).

artefact versus cognitive labour	ALARM CLOCK VERSUS KNOCKER-UPPER	CAMERA VERSUS HUMAN VISION	GARMENT FACTORY VERSUS SEAMSTRESS/TAILOR
			
			
a) label	replacement	enhancement	displacement
b) valence	neutral	beneficial	harmful
c) effect on cognition	reskilling	unaffected	deskilling
d) labour obfuscation	minimal	moderate	maximal
e) human equivalence	worse or same	different or better	quantity, not quality
f) human-in-the-loop	none	user	workers
g) human input	current time, ring time, energy source	framing, memory or film, energy source	pattern, materials, workers, wages
h) desired output	user awake	persistent image or video	clothes

The same holds for the human computers in the USA at the National Aeronautics and Space Administration (NASA; Shetterly, 2016b):

Early on, when they said they wanted the capsule to come down at a certain place, they were trying to compute when it should start. I said, “Let me do it. You tell me when you want it and where you want it to land, and I’ll do it backwards and tell you when to take off.” That was my forte. [...]

But when they went to computers, they called over and said, “tell her to check and see if the computer trajectory they had calculated was correct.” So I checked it and it was correct. (Katherine Johnson, as quoted in Hodges, 2008, n.p.)

And in both countries these women were written out of the historical record; their erasure facilitated by the rise of the electronic computer. This general pattern of displacement of women by machines, which we have previously dubbed *Pygmalion displacement* (Erscoi et al., 2023), and that of people by machines generally, has had harmful effects on society (Adler, 1990; Agar, 2003; Sherwood, 1985; Wendling, 2002), like the permanent harm to the British computer industry in the 20th century (M. Hicks, 2017). Importantly, “digital computers were

promoted as more efficient and less error-prone than humans at calculations. But in fact this comparison is not 'like-for-like' since, for example, calculating ballistic trajectories is, when done by women, also open to ethical questioning." (Erscoi et al., 2023, p. 20)

Centring the human cognitive component, as well as outlining the artefact in the relationship with the definition in Table 1, teases out important differences between the presented relationships. For example, all three pairs have the same desired output, which is the result of the calculation (row *h*, Table 2) the abacus requires all of human cognition to work, (rows *f* and *g*), while the calculator really does take over arithmetic, and the digital computer can take over even more, assuming the programmer can code it: "algorithms are always already made, maintained, and sustained by humans." (Bucher, 2018, p. 52) This also underlines how easily significant cognitive labour can be obfuscated when we move from left to right in Table 2: every use of the abacus has obvious manipulation effort while once a computer is programmed, software runs without any direct indication it was handcrafted (row *d*). In other words and relevant for our fields, "given the Cartesian legacy of the cognitive sciences, computers are looked at with veneration as soon as they produce well-formed output (Weizenbaum, 1976)" (Rasenberg et al. 2023, p. 312; Jucan 2023; Powell 1970). We should therefore be on high alert when others (or we) program computers to perform complex tasks, so as not to be taken in by this and misled into thinking something mystical — something other than a machine obeying our formal instructions — has occurred.

### 3.2 Alarm Clock, Camera, Garment Factory

Moving away from arithmetical operations and Turing-complete comparisons, to specific artefacts outside obviously computational devices: Table 3 depicts the AI relationships between alarm clocks, cameras, and garment factories and respective selected cognitive capacities. An alarm clock is a simple device that we provide some basic inputs to (row *g*, column Table 3) for it to function as we desire, e.g. to ring at a specific time (row *h*). To the untrained eye, an alarm clock may appear non-cognitive, and yet depicted in the black & white photo in column 1, Table 3: "Mrs Mary Smith wakes the dockers of Limehouse, London, with her peashooter in 1931. [...] She was a knocker-upper, a human alarm clock" (Topham, 1931, n.p.). The alarm clock completely automates every aspect of her kind of profession, with the exception there is no guarantee the user will be awake (row *e*, column Table 3). Unlike with a human alarm clock, an artefact cannot promise the desired output — we all have experienced sleeping through loud noises or some other alarm malfunction, and on the flip-side children are often woken by their parents to ensure they make it to school. Additionally, the inclusion of alarm clock functionality in mobile phones means their use can be further generalised through the day as reminders, with the inputs staying the same as an old-fashioned alarm clock and the desired output always requiring human supervision, e.g. a reminder to take out the rubbish is merely a reminder and not a guarantee the bags are taken out. All this is very familiar to those experiencing struggles with the cognitive capacity of executive function.

A slight tangent here is useful on the history of clocks, which is one of control of the users by the measurement of time, and not the other way round, which reflects and underlines the need to reorient all AI into human-centric focus. Prior to clocks, human labour was governed by the natural passing of time sans measurement, e.g. waking with the sun. Clocks are hegemonic tools, backbones of industrialism and capitalism:

The problem of the clock is, in general, similar to that of the machine. Mechanical time is valuable as a means of co-ordination of activities in a highly developed society, just as the machine is valuable as a means of reducing unnecessary labour to the minimum. Both are valuable for the contribution they make to the smooth running of society, and should be used insofar as they assist [people] to co-operate efficiently and to eliminate monotonous toil and social confusion. But neither should be allowed to dominate [people's] lives as they do today. (Woodcock, 1944, p. 8)

This aspect of machines, when they measure and control us, is one to bear in mind, and one which we will return to time and again below.

The camera is perhaps a more understandable addition to Table 3 with its relationship to the human capacity of vision well-known. The point here — as with all the other relationships — is not a mechanistic similarity but one based on functional correspondence or role, factual or perceived (cf. Guest & Martin, 2023, 2024; Guest, Scharfenberg, & van Rooij, 2025). Notable in our relationship with the camera is the fact it has largely deskilled nothing in the present



for the typical user, using a camera or looking at its output, does not negatively impact our ability to visually perceive (rows *b* & *c*, Table 3). Photographs also serve, especially in the advent of mobile phone cameras, to help us recollect our memories, enhancing our own ability to think about the past (row *a*). Of course, the user, the photographer, must frame and perform automatic or manual adjustments to the lens, for any image or video to be of use (row *g*). However, this human contribution, which manifests as the user, is ever present in all our AI systems from calculators to alarm clocks to much more complex systems like digital computers. The more infrequently discussed human-in-the-loop is humans other than the user, such as cases like digital computers (3rd column of Table 2) which require many other people's labour, not just the user's (Birhane, Han, et al., 2023; Birhane, Prabhu, et al., 2023; Birhane et al., 2024; Couldry & Mejias, 2019; Crawford, 2021; Kalluri et al., 2023, 2025; Placani, 2024). An infamous example of the human-in-the-loop technique is the orientalist Mechanical Turk, which toured from the late 18th to the mid 19th century, wherein a person hid in a cabinet under what appeared to be an automaton that played chess (Stephens, 2023). In fact, the person below the puppet controlled its movements, giving the human player sitting across it the impression that they were being beaten at chess by a clockwork machine. This is also the namesake of Amazon's Mechanical Turk, a platform on which low paid workers toil to "earn pennies or dollars doing tasks that computers cannot yet [and may never] easily do." (Newman, 2019, n.p.)

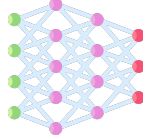





The sociotechnical relationship between the garment factory, which comprises humans-in-the-loop, can be seen as a more harmful, and equally obfuscatory, variant of the Mechanical Turk. Workers are treated badly in many cases, such as where sweatshop labour is used, the environmental impact of so many clothes is ignored, and the harmful chemicals are glossed over. In the not so distant past, people owned fewer clothes and had them hand- and custom-made by seamstresses and tailors. We readily have accepted this displacement of a relationship that involved a skilled adult making us a small selection of well-fitting clothes, which typically could last a whole lifetime, to a world in which we consume clothes relentlessly made often by workers, who can even be underage, and are often in harmful conditions (cf. Crawford, 2021; Nick Dyer-Witheford, 2019; Perrigo, 2023; Steinhoff, 2021; Stephens, 2023; Wendling, 2002). Notwithstanding, bespoke tailoring is to this day understood to be superior, because it factually is, and so preferred by the rich and famous and indeed required by anybody outside the bounds of factory-made standardised clothing configurations and sizes (row *e*). Even the simpler skill (compared to sewing from scratch) of taking in or out clothes as our bodies change over time is abandoned (row *c*) as a function of the garment factory which produces new clothes cheaply that we can buy instead (row *e*). We now turn to the cutting edge of contemporary AI.

### 3.3 LLM, Image Generator, Chatbot

If we centre the cognitive labour when teasing apart these three AI relations, as Table 1 guides us to do, it falls into our lap how it is displacement, harmful, and deskilling in every case: large language model (LLMs) versus essay writing, image generator versus artist, and chatbot versus human companionship (rows *a*–*c*, Table 4). Importantly, as Rasenberg et al. (2023, p. 312) explain, "whereas in human-animal interaction there is ample evidence of reciprocal adaptation, here the adaptation is strikingly one-sided, with [such systems] essentially helpless, requiring care (Lipp, 2023) and forcing people to adapt to their constraints (Alač et al., 2020; Suchman, 2019)." We are expected somehow to learn what is presented as a non-skill, because 'prompt engineering' is indeed not a skill, to enable us to use these opaque corporate-owned stochastic context-addressable systems. And, as turbo charged versions of the Mechanical Turk, contemporary AI models "are built on massive exploitative 'ghost labour'; crowdsourced and outsourced labour that follows the patterns of colonial relations (e.g., Bender et al., 2021)" (McQuillan et al., 2024, p. 3).

The false advertising — there is no engineering in prompt engineering — grows the longer the rhetoric around these systems are examined because, as we have seen many times so far, "data can only do so much. [In any apparently successful application of AI t]he real work is carried out by the people on the ground" (Dorrell, 2025, n.p.). While in the previous century, "computing systems functioned due to vast arrays of human workers, expressed through workflow organization, operators' actions, and software" (M. Hicks, 2017, p. 5), in the present the obfuscated human-in-the-loop — such as the sweatshop workers who guide and power LLMs and other such systems, or us ourselves whose data is stolen without our knowledge — is not respected (Birhane & Guest, 2021; Erscoi et al., 2023; Guest & Forbes,

Table 4. Contemporary AI products and models, such as LLMs, image generators, and chatbots are framed such that they are in competition with, or seen as equivalent to, cognitive capacities like essay writing, creating artwork, and providing companionship.

artefact versus cognitive labour	LLM VERSUS ESSAY WRITING	IMAGE GENERATOR VERSUS ARTIST	CHATBOT VERSUS COMPANIONSHIP
			
			
a) label	displacement	displacement	displacement
b) valence	harmful	harmful	harmful
c) effect on cognition	deskilling	deskilling	deskilling
d) labour obfuscation	maximal	maximal	maximal
e) human equivalence	worse	worse	worse
f) human-in-the-loop	data, programmers, sweatshop workers	data, programmers, sweatshop workers	data, programmers, sweatshop workers
g) input	so-called prompts	so-called prompts	so-called prompts
h) desired output	essay matching prompt	image matching prompt	unclear, wellness

2024; McQuillan, 2022; O'Neil, 2016). Even the low bar human computers' treatment has set is not met by modern AI's dehumanisation and theft of labour (e.g. Brennan et al., 2025; Crawford, 2021; Goetze, 2024; Perrigo, 2023; Rhee, 2018; van der Gun & Guest, 2024). In the time of software updates, the obfuscation of programmers' labour with seamless updates grows, while in the past one would physically go get artifacts fixed (or indeed fix them ourselves).

The systems in Table 4, embody an obfuscation of labour so complete the user believes the machine thinks for itself (Polo, 2024). In reality, exploited sweatshop workers in the Global South who perform the human-in-the-loop role do a lot of what we consider automated by AI (Bainbridge, 1983; Bender et al., 2021; Brennan et al., 2025; Crawford, 2021; McQuillan et al., 2024; Perrigo, 2023; Strauch, 2017).

There it is a definite social relation between [people], that assumes, in their eyes, the fantastic form of a relation between things. In order, therefore, to find an analogy, we must have recourse to the mist-enveloped regions of the religious world. In that world the productions of the human brain appear as independent beings endowed with life, and entering into relation both with one another and the human race. So it is in the world of commodities with the products of [people]'s hands. (Marx, 1867)

The reasoning problems become evermore severe in our misunderstandings of these most modern machines (Guest & Martin, 2023, 2024, 2025; Guest, Scharfenberg, & van Rooij, 2025). As Taina Bucher (2018, p. 50) explains: “When a machine runs smoothly, nobody pays much attention, and the actors and work required to make it run smoothly disappear from view (Latour, 1999).” A next step in this devolution and devaluation of cognitive labour is that now the user too, deskilled and displaced, also disappears from view. What voice does the human, now reduced to *only* a user, have if their essays (column 1, Table 4), their visual expression (column 2) are just the copy-pasted output of a device that performs patchwork plagiarism? What human connection do they have when their friends and romantic partners have been replaced with inanimate objects (column 3)?

The technopositive argument used to be that violence in video games was not indicative or causative of interpersonal violence because video game characters are virtual and users know the interaction is not in any way equivalent to that with other people outside the game. If that logic still holds then users will gain no, or very few, positive effects if they are lonely and need companionship (column 3, Table 4). If that logic does *not* hold, and people see these entities as possessing minds or as people, much more has unravelled. Indeed, for certain users, the psychological damage caused by only or mostly interacting with entities known to be designed to trick the user into believing they are people is immense (Dillon, 2020; Dupré, 2025; Erscoi et al., 2023; Hill, 2025; Jucan, 2023; Placani, 2024; Strengers et al., 2024; Turkle, 1984; Turkle et al., 2006; Weizenbaum, 1966, 1976; even the companies involved accept this potential harm: Phang et al., 2025).

As mentioned, past worries about sociotechnical relations were along these lines:

In a sane and free society such an arbitrary domination of [humanity’s] functions by either clock or machine would obviously be out of the question. The domination of [humanity] by the creation of [humanity] is even more ridiculous than the domination of [humanity] by [humanity]. (Woodcock, 1944, p. 8)

But in the present, and as Marx (1867) notes above, there is the domination of humans by other humans via religious-, conspiratorial-, or cult-like logic (Dupré, 2025; Guest & Martin, 2023, 2024; Hao, 2025; Heffernan, 2025; Reader, 2024; Samuel, 2023) and through these machines, algorithms, models. Humans’ expressions and social relations are not even mediated through technology, like when using a texting application to communicate with another person, but constitute technology as such. Technology, in this scheme, controlled by a private company (as all examples in Table 4) produces our so-called self-expressions. These relations are no longer captured by already flawed metaphors like the echo chamber — the echo has gone, the chamber is devoid of people: we neither shout nor are heard. We abandon our voice, forget how to use it, and forfeit what makes humans special in all columns of Table 4. In centring cognition as the analytical tool in Table 1 ensures, we are forced to look the human-in-the-loop in the eyes and recognise that these relations are harmful to people.

## 4 Machine Hauntology & Spectral Technology

Ignoring the ghosts in our machines harms us and our understanding of all such systems, from scientific models (Guest & Martin, 2021; Morgan & Morrison, 1999; Powell, 1970; van Rooij et al., 2024) to chatbots (Dillon, 2020; Erscoi et al., 2023; Jucan, 2023; Turkle et al., 2006). To truly centre the human in AI, we must admit the human’s direct and inherent centrality, and such an admission can be facilitated by the radical redefining of AI in Table 1 and the examples unpacked in Tables 2–4. And so if HCAI wishes to uphold and enact its human-centred-ness, as a field or perspective, it must then aspire to properly address each of the four points that are core to its current formulation (listed in section 1, repeated here), then it must:

1. with respect to supporting or enhancing human skills, both without displacement of said skills and without violation of fundamental human rights, recognise and **act when displacement AI relationships take place**;
2. with respect to imbuing systems with various so-called human-aligned values, including those of explainability and transparency, realise that **human-aligned values can only exist in systems where we actively uphold those values**, they do not come for free, and are not formally guaranteed, but are constantly negotiated through sociotechnical struggles;

3. with respect to focussing on human behaviour as a benchmark, i.e. the idea of human-like or -level performance, steer clear of correlationist, such as naive computationalist or modern connectionist, stances and roundly **reject benchmarks as meaningful**;
4. with respect to implicating behavioural, or otherwise psychological, methods in the study of these systems as if on equal epistemic footing with humans, take heed of serious warnings about our collective scientific reasoning as **correlations are red herrings in the search for theoretical understanding**.

To do this, as I have demonstrated, the definition proposed in Table 1 changes how we see artefacts, freeing us to view many more artefact-human relations as having human cognition at their centres. For the first part of my new proposed definition, if cognitive labour appears to be outsourced to a greater or lesser extent to an inanimate object, we can call this relationship between technology and cognition: AI. For the second step, even more analysis is needed wherein we need to discern and evaluate the relationship (recall rows in Table 1) between humans and a given use of the artefact under question (recall Tables 2–4). Deflating allows us to do things that chasing after what AI currently is per fads of the technology sector does not — specifically:

- a. We can **centre human cognition**, and therefore we can centre the study thereof, cognitive science, as a relevant discipline to understand purported cases of artificial cognition (e.g. van Rooij et al., 2024). Under computationalism, which the possibility of engineering cognition implicates, correlations and benchmarks are not relevant (e.g. Guest & Martin, 2025; Guest, Scharfenberg, & van Rooij, 2025). This allows us to in principle reject any argument that uses behavioural or neuroimaging correlation to argue for human-likeness of an artifact.
- b. We can **reject AI hype**, anthropomorphism, mysterianism about known mechanism,<sup>4</sup> exaggeration, or fads, removing this rhetoric from being relevant as to what counts as AI, which is a typical frame with many definitions of AI, often to push products (e.g. Duarte et al., 2024; Forbes & Guest, 2025). The same goes for claims about neural or biological plausibility or inspiration; these have no useful coherent definitions (e.g. Guest & Martin, 2023, 2024).
- c. We can **consider AI in general abstracted terms** without requiring specific reference to current advances in AI, e.g. chess-playing algorithms are AI, regardless of whether systems are cutting edge or not (sometimes called the AI effect; McCorduck, 2004). And so we can easily reject that *only* artificial neural nets or large language models or generative AI are AI, especially when until recently GOFAI (good old-fashioned AI, also known as symbolic) was canonically AI, hence the name (Guest, Suarez, et al., 2025).
- d. Relatedly, we can **grant AI a (pre)history**, allowing us to include the Antikythera mechanism (Freeth et al., 2021), astrolabes, sextants, abacuses, and more in the timeline of AI (e.g. Erscoi et al., 2023; Mayor, 2018). We can uncondense time — allowing us to slow down and giving us back our history — which is centrally relevant for understanding our present or possible futures (Hamilton, 1998; Stengers, 2018). As mentioned, this is something of a phobia, notably:

In English, the use of the word cybernetics raises no difficulties. Frenchmen with sufficient curiosity, however, were surprised to find it in Littré and Larousse; and the forgotten writings of Ampère were exhumed. When someone eventually turned up the new term in Plato, some of the experts rose in horror, declaring that kybernitiki should on no account be translated 'cybernetics'. (Guilbaud, 1960)

Taken together, these properties and by-products of Table 1 allow us to perform transcendental as well as immanent analyses of AI, such that we can pick out more than artificial neural networks, or specifically large language models, or such that we can perform analyses outside the tired contrasts of GOFAI or symbolic AI versus connectionist AI (Guest & Martin, 2024). No more mystification is possible because clarity and simplicity of definitions is within

<sup>4</sup>The idea that engineered mechanisms, like matrix multiplication or other operations in artificial neural networks are somehow uniquely unknown or unknowable properties of these systems.



reach, and because these models are now correctly positioned on a historical timeline, and subject to scientific investigation outside the correlationist dogma (Guest & Martin, 2024, 2025; Guest, Scharfenberg, & van Rooij, 2025). We do not need to passively be led astray, and in circles, by the technology sector any more, which “is mostly concerned with building profitable artifacts and is unconcerned with abstract definitions of intelligence.” (Heffernan, 2019, p. 4)

Instead we can focus on: finding persistent themes — e.g. Pygmalion displacement: the displacement of women by machines and algorithms, a process that involves dehumanising the women while humanising the technologies (Erscoi et al., 2023) — through teasing out properties of the sociotechnical relationship; avoiding hype as we centre human cognition; and using example artifacts from the distant and near past to help cut through exaggerated claims. Importantly, we can focus on the harms of the sociotechnical relationship on a case-by-case basis, while also learning from the past without getting bogged down by whether a specific system is, e.g. generative AI versus a convolutional neural network, which is discriminative AI (Efron, 1975; Guest, Suarez, et al., 2025; Jebara, 2004; Mitchell, 1997; Ng & Jordan, 2001; Xue & Titterton, 2008). We sidestep being tricked into using such formal terminology wrongly in service of the technology industry’s spin game, which coopts and distorts formal terms to cause confusion and hype (Duarte et al., 2024; Guest, Suarez, et al., 2025; Helfrich, 2024). This kind of protection from misinformation is especially important as companies largely deploy closed source models, but even in the case of open source code often provide misleading or otherwise lacking documentation, limiting scientific investigation (Barlas et al., 2021; Birhane, Prabhu, et al., 2023; Birhane et al., 2024; Jackson, 2024; Kalluri et al., 2025; Liesenfeld et al., 2023; Mirowski, 2023; Ojewale et al., 2025; Widder et al., 2024).

Relatedly, the cyclic reasoning of treating machines as if we have decided already that they are human-like (points 3 & 4 above) is evident at the birth of AI as a field. As Teresa Heffernan (2024), explains:

Turing speculated that by the end of the century the “use of words...will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted” (1950, p. 442), and many of today’s AI researchers have, following Turing’s lead, altered the meaning of words like — “reading,” “intuiting,” “feeling,” “dreaming,” and “creating” — to accommodate machine logic.

And the same for other so-called founding fathers of AI, who when coining ‘Artificial Intelligence’ claimed “that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it” (McCarthy et al., 1955, p. 2). This foundational document of AI, which was a proposal for a two month long summer conference, shows all the hallmarks we wish to avoid. As described by Pettit (2024, p. 47): “The starting point for this multidisciplinary effort was the coupling of ‘natural and artificial intelligence,’ although, as the Dartmouth proposal made clear, the distinction between the two was uninteresting for many. These two realms constituted a common problem.” In a now classic act of condensing time (recall Hamilton, 1998; Stengers, 2018), McCarthy et al. (1955) infamously proposed that this project of ‘solving’ intelligence would take a single season. Not in this AI summer, and not in any, will mathematically impossible displacements of humans pass muster (Rich et al., 2021; van Rooij et al., 2024).

More and more automation, much like a longer and longer road, does not imply that along the way something magical starts to happen merely as a by-product of distance between cause and effect or origin and destination. Nothing qualitatively different than creating a series of pathways between locations is underway if we keep extending the road. This being said, what does qualitatively shift is the a posteriori tractability of the emergent possible pathways if roads intersect and branch out. My specific route to work is tractable and manageable by me, but the minutiae of the route are likely impossible to guess without further information the longer the distance is as the road network, and my own preference for stops along the way, for example, together provide infinite possibilities for which route I *could* take. There is no magic here — other than metaphorically used perhaps for the awe human cognition should inspire — the proverbial ghost in the machine is the literal human-in-the-loop (Guest & Martin, 2025). The route is tractable to me because I picked it, while also being intractable to guess given a large enough search space for somebody else to copy.

This same metatheorising could be at play when we wrongly consider that engineered systems even more complex than roads as having qualitative shifts the more expansive the automation is — consider an abacus versus a calculator versus a full-blown digital computer



(as in Table 2). The only qualitative shift is in their formal complexity which is not a function of a quantitative aspect of a system. The digital computer is closer to literal Turing-completeness, but no slight of hand is at play: we know computers are made-up of smaller components such as flip-flops and logic gates and it is not their *quantity* that causes Turing-completeness. Amusingly, although perhaps also depressingly for the state of computational literacy, LLMs are most likely less computationally expressive than the digital computers on which they run. What is the case is the road, the conceptual and actual distance, from bare metal and silicon hardware to the software is longer. LLMs appear above the physical computer on which they run, seemingly ensnaring us to assume such an extension of matter touches on the cognitive. But like a longer road nothing has really changed in the engineered system: just more of the same. This is unlike the biological and cognitive spheres. And while unintuitive for many, the acceptance of this difference is the way forwards under computationalism.

I do not say the above to motivate rejections of mainstream computationalism, but if the naive conception furthers mystifying basic computational processes and engineering then it surely must, we must make it, give way to more mature ways of thought that take into account what it is we commit to (cf. basic versus naive versus non-naive computationalisms, Guest & Martin, 2025). In other words, just because the destination along the road is currently invisible to the naked eye from the origin does not by any stretch of the imagination imply something magical will get us there. What gets us to the destination — any goal — is our pre-existing cognition whether the end is currently visible to the naked eye or not. When it comes to the part of getting there that is automated, from car to calculator, the systems' parts are known and just hum along similarly regardless of whether the distance is 50 meters or 50 thousand or whether the digits are in the tens or the trillions.

Those who dogmatically insist that benchmarks have something to say about the humanity of machines assume that a behavioural match is informative, and that the correlationist programme can deliver evidence for their assumption (Guest & Martin, 2021, 2024). But cyclic reasoning collapses here under its own weight. There are infinite correlations to draw between an item and its reflection, but nothing in the, AI or traditional, mirror is the thing itself (Vallor, 2024). Such

dogmatists are lazy-bones. They refuse to undertake any painstaking study of concrete things, they regard general truths as emerging out of the void, they turn them into purely abstract unfathomable formulas, and thereby completely deny and reverse the normal sequence by which [humanity] comes to know truth. (Tse-Tung, 1937)

Truths, correlational, or otherwise do not emerge from the data. They are the products of our cognition, and of our interactions with the world. It is perhaps an uncomfortable truth that there are parts of cognition, like the human practice of science, that cannot be automated (Rich et al., 2021; Ryle, 1949; Tanney, 2009; van Rooij, 2008).

There are no Rube Goldberg-like perpetual motion-like human-in-the-loop-free machines. There *never* will be. We might crave such a machine, but the only entities we know that are self-sustaining, autopoietic, is everything *but* machines. The only person who can decide to turn a computer off and then on again is exactly that: a person (Polo, 2024). To effectively avoid correlationism, obfuscation of cognition in AI, we must reject the current mainstream view: no amount of high scores on benchmarks, or any other correlational evidence, can ever pile up high enough to graduate to a causal claim. To de-fetishise AI, we must accept that AI is any relationship between technology, tools, models, machines and humans where it appears as if some cognitive labour is offloaded onto such artifacts, and furthermore we must accept that such a relationship requires methodical teasing apart to obviate the centrality of the human (Bernardi, 2024; Braune, 2020; Morris, 2017; Mota & Cosentino Filho, 2024; Pfaffenberger, 1988). We cannot rid the machine from its ghost. But we can rid the concept of human from ghostliness — the human “need not be degraded to a machine by being denied to be a ghost in a machine.” (Ryle, 1949, p. 301)

## References

- Adler, P. S. (1990). Marx, Machines, and Skill. *Technology and Culture*, 31(4).  
 Agar, J. (2003). *The government machine: a revolutionary history of the computer*. MIT press.

- Alač, M., Gluzman, Y., Aflatoun, T., Bari, A., Jing, B., & Mozqueda, G. (2020). Talking to a toaster: how everyday interactions with digital voice assistants resist a return to the individual. *Evental Aesthetics*, 9(1), 3–53.
- Andersen, T. O., Nunes, F., Wilcox, L., Coiera, E., & Rogers, Y. (2023). Introduction to the special issue on human-centred AI in healthcare: Challenges appearing in the wild.
- Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779.
- Bannon, L. (2011). Reimagining HCI: toward a more human-centered perspective. *interactions*, 18(4), 50–57.
- Barlas, P., Kyriakou, K., Guest, O., Kleanthous, S., & Otterbacher, J. (2021). To "see" is to stereotype: Image tagging algorithms, gender recognition, and the accuracy-fairness trade-off. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW3), 1–31.
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, O., & Mariman, R. (2024). Generative AI can harm learning. Available at SSRN, 4895486.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623.
- Bernardi, G. (2024). Limites absolutos ou relativos para a substituição de postos de trabalho por inteligência artificial? *Revista Eletrônica Internacional de Economia Política da Informação da Comunicação e da Cultura*, 26(2), 109–124.
- Birhane, A., Dehdashtian, S., Prabhu, V., & Boddeti, V. (2024). The dark side of dataset scaling: Evaluating racial classification in multimodal models. *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1229–1244.
- Birhane, A., & Guest, O. (2021). Towards Decolonising Computational Sciences. *Kvinder, Køn & Forskning*, 29(1), 60–73.
- Birhane, A., Han, S., Boddeti, V., Luccioni, S., et al. (2023). Into the LAION's den: Investigating hate in multimodal datasets. *Advances in Neural Information Processing Systems*, 36, 21268–21284.
- Birhane, A., Prabhu, V., Han, S., & Boddeti, V. N. (2023). On hate scaling laws for data-swamps. *arXiv preprint arXiv:2306.13141*.
- Bishop, J. M. (2021). Artificial intelligence is stupid and causal reasoning will not fix it. *Frontiers in Psychology*, 11, 513474.
- Boden, M. A. (2006). *Mind As Machine: A History of Cognitive Science Two-Volume Set*. Oxford University Press, USA.
- Braune, S. (2020). Fetish-oriented ontology. *Open Philosophy*, 3(1), 298–313.
- Brennan, K., Kak, A., & West, S. M. (2025, June). *Artificial Power: AI Now 2025 Landscape*. AI Now Institute.
- Bucher, T. (2018). *If...Then: Algorithmic Power and Politics*. Oxford University Press, USA.
- Couldry, N., & Mejias, U. A. (2019). Data Colonialism: Rethinking Big Data's Relation to the Contemporary Subject. *Television & New Media*, 20(4), 336–349.
- Crawford, K. (2021). *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Dillon, S. (2020). The Eliza effect and its dangers: from demystification to gender critique. *Journal for Cultural Research*, 24(1), 1–15.
- Dorrell, T. (2025). Cybernetics with Chinese Characteristics: How big data is eliminating poverty and building socialism. *People's World*. <https://www.peoplesworld.org/article/cybernetics-with-chinese-characteristics-how-big-data-is-eliminating-poverty-and-building-socialism/>
- Dreyfus, H. L. (1965). Alchemy and artificial intelligence.
- Duarte, T., Barrow, N., Bakayeva, M., et al. (2024). Editorial: The ethical implications of AI hype. *AI Ethics*, 649–651.
- Dupré, M. H. (2025, June). People Are Being Involuntarily Committed, Jailed After Spiraling Into "ChatGPT Psychosis". <https://futurism.com/commitment-jail-chatgpt-psychosis>
- Efron, B. (1975). The efficiency of logistic regression compared to normal discriminant analysis. *Journal of the American Statistical Association*, 70(352), 892–898.
- Egan, F. (2025). *Deflating mental representation*. MIT Press.
- Elgin, C. Z. (2017). *True enough*. MIT press.
- El-Mhamdi, E.-M., Farhadkhani, S., Guerraoui, R., Gupta, N., Hoang, L.-N., Pinot, R., Rouault, S., & Stephan, J. (2022). On the impossible safety of large AI models. *arXiv preprint arXiv:2209.15259*.

- Erscoi, L., Kleinherenbrink, A. V., & Guest, O. (2023, February). Pygmalion Displacement: When Humanising AI Dehumanises Women.
- Forbes, S. H., & Guest, O. (2025). To improve literacy, improve equality in education, not large language models. *Cognitive Science*, 49(4), e70058.
- Freeth, T., Higgon, D., Dacanalis, A., MacDonald, L., Georgakopoulou, M., & Wojcik, A. (2021). A Model of the Cosmos in the ancient Greek Antikythera Mechanism. *Scientific reports*, 11(1), 5821.
- Geyer, W., Weisz, J., Pinhanez, C. S., & Daly, E. (2022, August). What is human-centered AI? <https://research.ibm.com/blog/what-is-human-centered-ai>
- Goetze, T. S. (2024). AI art is theft: Labour, extraction, and exploitation: Or, on the dangers of stochastic Pollocks. *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 186–196.
- Grier, D. A. (2013). *When computers were human*. Princeton University Press.
- Guattari, F., Sheed, R., Cooper, D. G. (G., & Guattari, F. (1984). *Molecular revolution: psychiatry and politics*. Penguin.
- Guest, O. (2024). What makes a good theory, and how do we make a theory good? *Computational Brain & Behavior*, 7(4), 508–522.
- Guest, O., & Forbes, S. H. (2024). Teaching coding inclusively: if this, then what? *Tijdschrift voor Genderstudies*, 27(2/3), 196–217.
- Guest, O., & Martin, A. E. (2021). How computational modeling can force theory building in psychological science. *Perspectives on Psychological Science*, 16(4), 789–802.
- Guest, O., & Martin, A. E. (2023). On logical inference over brains, behaviour, and artificial neural networks. *Computational Brain & Behavior*, 6(2), 213–227.
- Guest, O., & Martin, A. E. (2024, October). A Metatheory of Classical and Modern Connectionism.
- Guest, O., & Martin, A. E. (2025, February). Are Neurocognitive Representations ‘Small Cakes’? <https://philsci-archive.pitt.edu/24834/>
- Guest, O., Scharfenberg, N., & van Rooij, I. (2025, May). Modern Alchemy: Neurocognitive Reverse Engineering. <https://philsci-archive.pitt.edu/25289/>
- Guest, O., Suarez, M., Müller, B. C. N., van Meerkerk, E., Oude Groote Beverborg, A., de Haan, R., & van Rooij, I. (2025). Against the Uncritical Adoption of ‘AI’ Technologies in Academia. *Manuscript in Preparation*.
- Guilbaud, G. T. (1960). What is cybernetics?
- Haigh, T. (2023). There Was No ‘First AI Winter’. *Communications of the ACM*, 66(12), 35–39.
- Hamilton, S. N. (1998). Incomplete Determinism: A Discourse Analysis of Cybernetic Futurology in Early Cyberculture. *Journal of Communication Inquiry*, 22(2), 177–204.
- Hao, K. (2025). *Empire of AI: Dreams and Nightmares in Sam Altman’s OpenAI*. Penguin Random House.
- HCAIM Consortium. (2025). <https://humancentered-ai.eu/>
- Heffernan, T. (Ed.). (2019). *Cyborg Futures: Cross-disciplinary Perspectives On Artificial Intelligence And Robotics*. Palgrave Macmillan.
- Heffernan, T. (2024). The imitation game, the “child machine,” and the fathers of AI. *AI & SOCIETY*, 39(1), 353–357.
- Heffernan, T. (2025). The religious roots of AI and the rise of neo-feudalism. *Religion*, 1–14.
- Helfrich, G. (2024). The harms of terminology: why we should reject so-called “frontier AI”. *AI and Ethics*, 4(3), 699–705.
- Hicks, M. (2017). *Programmed inequality: How Britain discarded women technologists and lost its edge in computing*. MIT press.
- Hicks, M. T., Humphries, J., & Slater, J. (2024). ChatGPT is bullshit. *Ethics and Information Technology*, 26(2), 1–10.
- Hill, K. (2025). They Asked ChatGPT Questions. The Answers Sent Them Spiraling. *The New York Times*. <https://www.nytimes.com/2025/06/13/technology/chatgpt-ai-chatbots-conspiracies.html>
- Hodges, J. (2008). She was a computer when computers wore skirts. *NASA*, 4, 2018. <https://www.nasa.gov/centers-and-facilities/langley/she-was-a-computer-when-computers-wore-skirts/>
- Holzinger, A. (2025). *About the Human-Centered AI Lab*. Retrieved April 13, 2025, from <https://human-centered.ai/about-the-holzinger-group/>
- Hughes, B. B. (1985). *World futures: A critical analysis of alternatives*.

- Human-Centered AI track, MSc Human-Technology Interaction, Eindhoven University of Technology. (2025). Retrieved April 13, 2025, from <https://www.tue.nl/en/education/graduate-school/masters-track-human-centered-ai>
- Human-centered Artificial Intelligence, Utrecht University. (2025, April). <https://www.uu.nl/en/research/human-centered-artificial-intelligence>
- Human-Centred AI theme, MSc Data Science and Artificial Intelligence Technology, TU Delft. (2025). <https://www.tudelft.nl/en/education/programmes/masters/dsait/msc-data-science-and-artificial-intelligence-technology/programme/themes/human-centred-ai>
- Jackson, S. (2024). Sam Altman Explains OpenAI's Shift to Closed AI Models. <https://www.businessinsider.com/sam-altman-why-openai-closed-source-ai-models-2024-11>
- Jebara, T. (2004). Generative versus discriminative learning. In *Machine learning: Discriminative and generative* (pp. 17–60). Springer.
- Jucan, I. B. (2023). *Malicious Deceivers: Thinking Machines and Performative Objects*. Stanford University Press.
- Kalluri, P. R., Agnew, W., Cheng, M., Owens, K., Soldaini, L., & Birhane, A. (2023). The surveillance AI pipeline. *arXiv preprint arXiv:2309.15084*.
- Kalluri, P. R., Agnew, W., Cheng, M., Owens, K., Soldaini, L., & Birhane, A. (2025). Computer-vision research powers surveillance technology. *Nature*, 1–7.
- Latour, B. (1999). *Pandora's hope: Essays on the reality of science studies*. Harvard university press.
- LePage-Richer, T. (2024). *Neural Networks*. meson press.
- Liesenfeld, A., Lopez, A., & Dingemanse, M. (2023). Opening up ChatGPT: Tracking openness, transparency, and accountability in instruction-tuned text generators. *Proceedings of the 5th international conference on conversational user interfaces*, 1–6.
- Lighthill, J., et al. (1973). Artificial intelligence: a paper symposium. *Science Research Council, London*.
- Lima-Silva, T. B., Barbosa, M. E. d. C., Zumkeller, M. G., Verga, C. E. R., Prata, P. L., Cardoso, N. P., Moraes, L. C. d., & Brucki, S. M. D. (2021). Cognitive training using the abacus: a literature review study on the benefits for different age groups. *Dementia & neuropsychologia*, 15, 256–266.
- Lipp, B. (2023). Caring for robots: How care comes to matter in human-machine interfacing. *Social Studies of Science*, 53(5), 660–685.
- Lu, Y., Li, M., Cui, Z., Wang, L., Hu, Y., & Zhou, X. (2023). Transfer effects of abacus training on cognition. *Current Psychology*, 42(8), 6271–6286.
- Marx, K. (1867). The fetishism of commodities and the secret thereof. Frederick Engels, trans. Samuel Moore; Edward Aveling. New York. <https://www.marxists.org/archive/marx/works/1867-c1/ch01.htm#S4>
- Marx Memorial Library. (2024). Is 'artificial intelligence' a threat or a promise? *Morning Star*. <https://morningstaronline.co.uk/article/%E2%80%98artificial-intelligence%E2%80%99-threat-or-promise>
- Mauss, M. (1923). Essai sur le don. Forme et raison de l'échange dans les sociétés archaïques. *L'Année Sociologique*, 30–180.
- Mayor, A. (2018). *Gods and robots: Myths, machines, and ancient dreams of technology*. Princeton University Press.
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (1955). A proposal for the Dartmouth summer research project on artificial intelligence. <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>
- McCorduck, P. (2004). *Machines who think: A personal inquiry into the history and prospects of artificial intelligence*. AK Peters/CRC Press.
- McQuillan, D. (2022). *Resisting AI: An anti-fascist approach to artificial intelligence*. Policy Press.
- McQuillan, D., Jarke, J., & Pargman, T. C. (2024). We are at an extreme point where we have to go all in on what we really believe education should be about. *Postdigital Science and Education*, 6(1), 360–368.
- Mirowski, P. (2023). The evolution of platform science. *Social Research: An International Quarterly*, 90(4), 725–755.
- Mitchell, T. M. (1997). *Machine learning* (Vol. 1). McGraw-hill New York.
- Moore, J. (1997). *A primitivist primer*. Green Anarchy Distro.
- Morgan, M. S., & Morrison, M. (1999). *Models as mediators: Perspectives on natural and social science*. Cambridge University Press.



- Morris, R. C. (2017). After de Brosses: Fetishism, translation, comparativism, critique. *The returns of fetishism: Charles de Brosses and the afterlives of an idea*, 133–319.
- Morrison, M., & Morgan, M. S. (1999). Models as mediating instruments. In M. S. Morgan & M. Morrison (Eds.), *Models as Mediators: Perspectives on Natural and Social Science* (pp. 10–37). Cambridge University Press.
- Mota, F. M. R., & Cosentino Filho, C. B. (2024). Fetichismo tecnológico no capitalismo de plataforma: relações e contradições da tecnologia e da IA com o Direito do Trabalho. *Revista Eletrônica Internacional de Economia Política da Informação da Comunicação e da Cultura*, 26(2), 81–108.
- Nado, J. (2021). Conceptual engineering via experimental philosophy. *Inquiry*, 64(1-2), 76–96.
- Newman, A. (2019). I found work on an Amazon website. I made 97 cents an hour. *The New York Times*, 15(2019).
- Ng, A., & Jordan, M. (2001). On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes. *Advances in neural information processing systems*, 14.
- Nick Dyer-Witthford, J. S., Atle Mikkola Kjösen. (2019). *Inhuman Power: Artificial Intelligence And The Future Of Capitalism*. Pluto Press.
- Novack, G. (1968). *Empiricism and Its Evolution: A Marxist View*. Pathfinder Press.
- Ojewale, V., Steed, R., Vecchione, B., Birhane, A., & Raji, I. D. (2025). Towards AI accountability infrastructure: Gaps and opportunities in AI audit tooling. *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–29.
- O’Neil, C. (2016). *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown.
- Pasquinelli, M. (2017). Machines that morph logic: Neural networks and the distorted automation of intelligence as statistical inference. *Glass Bead*, 1(1), 1.
- Perrigo, B. (2023). Exclusive: OpenAI used Kenyan workers on less than \$2 per hour to make ChatGPT less toxic. *Time Magazine*, 18, 2023. <https://time.com/6247678/openai-c-hatgpt-kenya-workers/>
- Pettit, M. (2024). *Governed By Affect: Hot Cognition and the End of Cold War Psychology*. Oxford University Press.
- Pfaffenberger, B. (1988). Fetishised objects and humanised nature: towards an anthropology of technology. *Man*, 236–252.
- Phang, J., Lampe, M., Ahmad, L., Agarwal, S., Fang, C. M., Liu, A. R., Danry, V., Lee, E., Chan, S. W., Pataranutaporn, P., et al. (2025). Investigating affective use and emotional well-being on ChatGPT. *arXiv preprint arXiv:2504.03888*.
- Pirozzoli, A. (2024). The human-centric perspective in the regulation of artificial intelligence. *European Papers-A Journal on Law and Integration*, 2024(1), 105–116.
- Placani, A. (2024). Anthropomorphism in AI: hype and fallacy. *AI and Ethics*, 4(3), 691–698.
- Polo, M. (2024). O pensamento de que a máquina é pensante: a necessária retomada da consciência crítica diante das invenções tecnológicas da humanidade. *Revista Eletrônica Internacional de Economia Política da Informação da Comunicação e da Cultura*, 26(2), 70–80.
- Ponomareva, V. (1998). The Human Factor in Space Exploration: Soviet and American Approaches. *Proceedings of the 1998 annual conference of the Institute for the History of Natural Science and Technology*, Moscow, 614–18.
- Powell, B. (1970). Descartes’ Machines. *Proceedings of the Aristotelian Society*, 71, 209–222.
- Raji, I. D., Kumar, I. E., Horowitz, A., & Selbst, A. (2022). The fallacy of AI functionality. *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 959–972.
- Raley, R., & Rhee, J. (2023). Critical AI: A field in formation. *American Literature*, 95(2), 185–204.
- Rasenberg, M., Amha, A., Coler, M., Van Koppen, M., van Miltenburg, E., De Rijk, L., Stommel, W., & Dingemanse, M. (2023). Reimagining language: Towards a better understanding of language by including our interactions with non-humans. *Linguistics in the Netherlands*, 40(1), 309–317.
- Reader, T. M. P. (2024). Silicon Valley’s Obsession With AI Looks a Lot Like Religion. *The MIT Press Reader*. <https://thereader.mitpress.mit.edu/silicon-valleys-obsession-with-ai-looks-a-lot-like-religion/>
- Rhee, J. (2018, March). *The robotic imaginary*. University of Minnesota Press.
- Rich, P., de Haan, R., Wareham, T., & van Rooij, I. (2021). How hard is cognitive science? *Proceedings of the annual meeting of the cognitive science society*, 43(43).



- Rogers, Y. (2022). Commentary: human-centred AI: the new zeitgeist. *Human-computer interaction*, 37(3), 254–255.
- Ryan, M. (2024). We're only human after all: a critique of human-centred AI. *AI & society*, 1–17.
- Ryle, G. (1949). *The Concept of Mind*. Routledge.
- Samuel, S. (2023, July). Why Silicon Valley AI prophecies just feel like repackaged religion. <https://www.vox.com/the-highlight/23779413/silicon-valleys-ai-religion-transhumanism-longtermism-ea>
- Saparina, E. V. (1966). *Cybernetics Within Us*. Peace Publishers.
- Schmager, S., Pappas, I. O., & Vassilakopoulou, P. (2025). Understanding Human-Centred AI: a review of its defining elements and a research agenda. *Behaviour & Information Technology*, 1–40.
- Schrimpf, M., Kubilius, J., Lee, M. J., Murty, N. A. R., Ajemian, R., & DiCarlo, J. J. (2020). Integrative Benchmarking to Advance Neurally Mechanistic Models of Human Intelligence. *Neuron*. [https://www.cell.com/neuron/fulltext/S0896-6273\(20\)30605-X](https://www.cell.com/neuron/fulltext/S0896-6273(20)30605-X)
- Sherwood, J. M. (1985). Engels, Marx, Malthus, and the Machine. *The American Historical Review*, 90(4).
- Shetterly, M. L. (2016a). Dorothy Vaughan. NASA. <https://www.nasa.gov/people/dorothy-vaughan/>
- Shetterly, M. L. (2016b). *Hidden figures: The story of the African-American women who helped win the space race*. William Morrow.
- Shneiderman, B. (2022). *Human-centered AI*. Oxford University Press.
- Smith, P., & Smith, L. (2024). This season's artificial intelligence (AI): is today's AI really that different from the AI of the past? Some reflections and thoughts. *AI and Ethics*, 4(3), 665–668.
- Stanford Institute for Human-Centered AI. (2019). Stanford HAI. <https://hai.stanford.edu/>
- Steinhoff, J. (2021). Automation and Autonomy. *SpringerLink*.
- Stengers, I. (2018). *Another Science is Possible: A Manifesto for Slow Science*. John Wiley & Sons.
- Stephens, E. (2023). The mechanical Turk: A short history of 'artificial artificial intelligence'. *Cultural Studies*, 37(1), 65–87.
- Strauch, B. (2017). Ironies of automation: Still unresolved after all these years. *IEEE Transactions on Human-Machine Systems*, 48(5), 419–433.
- Strengers, Y., Phan, T., Duque, M., & Dahlgren, K. (2024). Anthropomorphizing Voice Assistants: A Research Agenda for Human-AI Relationships. In L. Fortunati & A. Edwards (Eds.), *The De Gruyter Handbook of Robots in Society and Culture* (pp. 391–410). De Gruyter.
- Suchman, L. (2019). Demystifying the intelligent machine. In *Cyborg futures: Cross-disciplinary perspectives on artificial intelligence and robotics* (pp. 35–61). Springer.
- Tambiana, M. (2019). EU guidelines on ethics in artificial intelligence: Context and implementation. *European Parliamentary Research Service*.
- Tanney, J. (2009). Ryle's Regress and Cognitive Science. In B. Ambrose & S. Laugier (Eds.), *La Philosophie d'Oxford au 20ème Siècle: Approches du Sens Commun*. Hildesheim: Olms. <https://kar.kent.ac.uk/10606/>
- Topham, J. (1931). EU058004. <https://www.topfoto.co.uk/asset/617>
- Tse-Tung, M. (1937). On contradiction. [https://www.marxists.org/reference/archive/mao/selected-works/volume-1/mswv1\\_17.htm](https://www.marxists.org/reference/archive/mao/selected-works/volume-1/mswv1_17.htm)
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460.
- Turkle, S. (1984). *The second self: Computers and the human spirit*. Mit Press.
- Turkle, S., Taggart, W., Kidd, C. D., & Dasté, O. (2006). Relational artifacts with children and elders: the complexities of cybercompanionship. *Connection Science*, 18(4), 347–361.
- Vallor, S. (2024). *The AI mirror: How to reclaim our humanity in an age of machine thinking*. Oxford University Press.
- van der Gun, L., & Guest, O. (2024). Artificial Intelligence: Panacea or Non-Intentional Dehumanisation? *Journal of Human-Technology Relations*, 2.
- van Rooij, I. (2008). The tractable cognition thesis. *Cognitive science*, 32(6), 939–984.
- van Rooij, I., Guest, O., Adolphi, F., de Haan, R., Kolokolova, A., & Rich, P. (2024). Reclaiming AI as a theoretical tool for cognitive science. *Computational Brain & Behavior*, 7(4), 616–636.

- Walton, N., & Nayak, B. S. (2021). Rethinking of Marxist perspectives on big data, artificial intelligence (AI) and capitalist economic development. *Technological Forecasting and Social Change*, 166, 120576.
- Wang, C. (2020). A review of the effects of abacus training on cognitive functions and neural systems in humans. *Frontiers in neuroscience*, 14, 913.
- Weizenbaum, J. (1966). ELIZA—a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), 36–45.
- Weizenbaum, J. (1976). Computer power and human reason: From judgment to calculation.
- Wendling, A. (2002). Partial Liberations. *International Studies in Philosophy*, 34(2), 169–185.
- Widder, D. G., Whittaker, M., & West, S. M. (2024). Why ‘open’AI systems are actually closed, and why this matters. *Nature*, 635(8040), 827–833.
- Wiener, N. (1948, October). *Cybernetics or Control and Communication in the Animal and the Machine*. The MIT Press.
- Wiener, N. (1950). *The Human Use of Human Beings*. Houghton Mifflin.
- Wills, A. J. (2025). *Special Issue: Advanced Studies in Human-Centred AI*. Retrieved March 15, 2025, from [https://www.mdpi.com/journal/behavsci/special\\_issues/657ATT8P1Q](https://www.mdpi.com/journal/behavsci/special_issues/657ATT8P1Q)
- Winther, R. G. (2021). The Structure of Scientific Theories. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2021). Metaphysics Research Lab, Stanford University.
- Woodcock, G. (1944). *The tyranny of the clock*. Wooden Shoe Books.
- Xie, Y., Chang, H., Zhang, Y., Wang, C., Zhang, Y., Chen, L., Geng, F., Ku, Y., Menon, V., & Chen, F. (2024). Long-term abacus training gains in children are predicted by medial temporal lobe anatomy and circuitry. *Developmental Science*, 27(4), e13489.
- Xue, J. H., & Titterton, D. M. (2008). Comment on “On discriminative vs generative classifiers: A comparison of logistic regression and naive bayes”. *Neural processing letters*, 28, 169–187.
- Yuval-Davis, N. (2024). Antisemitism is a Form of Racism—or is it? *Sociology*, 58(4), 779–795.