Inhuman Power

Digital Barricades:
Interventions in Digital Culture and Politics

Also available:

Shooting a Revolution:
Visual Media and Warfare in Syria
Donatella Della Ratta

Cyber-Proletariat:
Global Labour in the Digital Vortex
Nick Dyer-Witheford

The Digital Party:
Political Organisation and Online Democracy
Paolo Gerbaudo

Gadget Consciousness:
Collective Thought, Will and Action in the Age of Social Media
Joss Hands

Information Politics:
Liberation and Exploitation in the Digital Society
Tim Jordan

Sad by Design:
On Platform Nihilism
Geert Lovink

Unreal Objects:
Digital Materialities, Technoscientific Projects and Political Realities
Kate O'Riordan

Inhuman Power

Artificial Intelligence and the Future of Capitalism

Nick Dyer-Witheford, Atle Mikkola Kjøsen
and James Steinhoff

PLUTO PRESS

First published 2019 by Pluto Press
345 Archway Road, London N6 5AA

www.plutobooks.com

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

ISBN 978 0 7453 3861 3 Hardback
ISBN 978 0 7453 3860 6 Paperback
ISBN 978 1 7868 0395 5 PDF eBook
ISBN 978 1 7868 0397 9 Kindle eBook
ISBN 978 1 7868 0396 2 EPUB eBook

This book is printed on paper suitable for recycling and made from fully managed and sustained forest sources. Logging, pulping and manufacturing processes are expected to conform to the environmental standards of the country of origin.

Typeset by Stanford DTP Services, Northampton, England

Contents

Introduction: AI-Capital

THE MOST VALUABLE THING

Capitalism is today possessed by the Artificial Intelligence (AI) question. Consider the Vancouver start-up Sanctuary Cognitive Systems Corporation, which aims to develop 'humanoid robots that can move, speak and think for themselves and interact – as intellectual peers – with real people'. Its owner, Geordie Rose, a quantum computing pioneer, concedes to an interviewer that, while building blocks towards this goal are already in common use in multiplying types of 'narrow AI' and specialized robotics applications, none are remotely close to an 'Artificial General Intelligence' (AGI) capable of full human emulation. His company's mission 'to unlock how human intelligence works and to replicate it on a mass scale' therefore 'sounds like a mind-boggling moonshot'. However, Rose is undeterred, for if he succeeds, it will be 'the most valuable thing ever created. What we're talking about is fundamentally altering the basis of capitalism itself' (Silicoff 2018).

It would be easy to dismiss this quest, were Sanctuary not competing against some of the most powerful capitalists in the world, also striving to produce AGI and related technologies: Elon Musk and his non-profit OpenAI; Vicarious FPC, Inc., backed by Samsung and technology billionaires Mark Zuckerberg and Jeff Bezos; and DeepMind Technologies, acquired by Google/Alphabet, whose owners, Sergei Brin and Larry Page, are patrons of the transhumanist Raymond Kurzweil (2005a), the most famous prophet of a 'technological singularity' in which computers attain human-equivalent intelligence. Sanctuary Cognitive Systems Corp may or may not survive (it is clearly cash-strapped and looking for angel investors). But its story of soaring technological ambition, mission-driven digital entrepreneurship and creepy androids 'like underworld creatures from a Hieronymus Bosch painting' (Silicoff 2018) is symptomatic of the AI-fever sweeping the world market, a fever that also manifests in a burgeoning business literature on AI applications, torrents of conflicting predictions about AI's consequences for employment, utopian speculation on the creation of 'Life 3.0' (Tegmark 2017), and fictions

ranging from pulp robo-apocalypses (Wilson 2012, 2015) to complex literary explorations on the new techno-existential horizon posited by AI (Mason 2017).

Defining AI is difficult. Nonetheless, Rose is correct that what is often termed 'narrow AI' is already present in the algorithmic processes that now inform much of everyday life. For warehouse workers or military personnel, such AIs may incarnate in the chassis of a robot delivery vehicle or semi-autonomous killer drone. Most AIs, however, act invisibly in the background of activities conducted on smartphones and computers; in search engine results, social media feeds, video games and targeted advertisements; in the acceptance or rejection of applications for bank loans or welfare assistance; in a call centre inquiry or summons to an on-demand cab; or in encounters with police or border guards, scanning their shadowed screens. In these ways, AI has been with us for years.

Once upon a time, people on the left referred to the regimes of the USSR and Eastern Europe as 'actually-existing socialism' (Bahro 1978), indicating an incipient but imperfect realization of hopes for a new social order. We propose an analogous formulation: 'actually-existing AI-capitalism', designating a phase of experimental and uneven adoption of the technologies in which so many hopes are invested. This phase may be protracted far longer than AI enthusiasts anticipate. It may stagnate, stall out and implode (as 'actually-existing socialism' did). But it could also intensify or expand in a transition either to a significantly transformed capitalism, or to a radically different social formation.

This reference to the fate of socialism brings us to the vantage point from which we critique AI-capitalism. Prophets of a technological singularity expect its arrival around 2045 (Kurzweil 2005a). This would put it 201 years after a critical observer of early industrial capitalism penned his own prediction as to its eventual outcome: 'finally … an inhuman power rules over everything' (Marx 1975 [1844]: 366). The young Marx was not writing about artificial intelligence or robots. He was describing the 'alienation' of workers dispossessed by capital of control over what they made, how they made it, their relations with fellow human beings and of their very 'species-being'. The argument we make in this book is that AI should be seen as the culmination of this process, a moment where the market-system assumes a life of its own. AI, we posit, is 'alien power' (Marx 1990: 716) – the power of autonomous capital. We read AI and Marxism through one another: AI through Marx, because Marx's

analysis of capitalism is the most comprehensive critical account of the fusion of commodification and technology driving AI forward today; Marx in the light of AI, because AI problematizes human exceptionalism, agency and labour in ways that profoundly challenge Marxist assumptions, and hence requires careful examination by those who share Marx's aspiration for revolution against and beyond capital.

THREE POLEMICS

The argument of this book interweaves three polemic critiques. The first is a critique of AI as an instrument of capital, with all this entails in terms of both the exploitation in and ejection from waged work of human labour, and the concentration of wealth and social power in the hands of the corporate owners of high technology. Depictions of AI as the outcome of a disinterested process of scientific research are naive. Machine intelligence is the product not just of a technological logic, but simultaneously of a social logic, the logic of producing surplus-value. Capitalism is the fusion of these technological and social logics and AI is the most recent manifestation of its chimerical merging of computation with commodification. Jump-started by the digital experiments of the US military-industrial complex, AI emerged and developed within a socio-economic order that rewards those who own the means for automating human labour, accelerating sales, elaborating financial speculation and intensifying military-police control over potential restive populations.

Whether or not AI may be put to different uses – 'reconfigured' (Bernes 2013; Toscano 2014; Steinhoff 2017) to contribute to or create a different social order – is a question we discuss later. What is apparent is that the owners of the great digital corporations regard AI as their technology – and with good reason, for it is they who possess the intellectual property rights, the vast research budgets, the labour-time of AI scientists, the data and the centres that store it, telecommunications networks, and the ties to an enabling state apparatus that are the preconditions for the creation of AI. It is they and their high-ranking managerial cadres who are in a position to implant their goals and priorities within AI software and hardware, 'baking-in' their values – in practice, the one prime directive, to expand surplus-value – to its design.

There may appear to be surprising diversity of opinion about AI amongst corporate leaders, ranging from ecstatic embrace to apocalyptic

warning. What is shared, however, is the tacit agreement that it is they who are to dictate the direction of AI, to determine in their high-level conclaves and privileged conversations with government how (not if) it is to be adopted, and with what balance between wild gladiatorial free-enterprise competition and cautionary ethical regulation and policy safety-nets to prevent unwelcome tumults. Whether it is Sergey Brin endorsing the idea of the singularity while consolidating his company's monopolistic powers to direct it, or Elon Musk warning of AI catastrophe while building (or attempting to build) fully automated factories for the production of self-driving vehicles, or Bill Gates offering feeble robot tax plans (and thereby drawing the instantaneous ridicule of peers), the colourful clashes of corporate personalities cover the more sober reality that these great AI moguls are no more, or less, than the personifications of abstract forces of market calculation that drive towards the maximization of profit. They also obscure the massive hubris of the capitalist class that believes it can control the forces it has unleashed. For we did not quite complete our titular quote from the young Marx: 'finally – *and this goes for the capitalists too* – an inhuman power rules over everything (1975 [1844]: 366, emphasis added).

In making this critique of capitalism's encounter with AI we also, however, take issue with leftist theorists who share such concerns, some of whom, like us, quote Marx in their evaluations of AI. So here we quarrel with interlocutors we respect and have learned from, but with whom we differ. There are two specific left perspectives against which we argue, perspectives that we dub 'minimizing' and 'maximizing' views on AI.

The left 'minimalist' position dismisses current discourses on AI as hype and hucksterism. A more moderate version grants them a limited credibility but insists this is not sufficient to seriously change previous analyses of capital and class (Huws 2014; Moody 2018a). For a strong statement of this minimalist position, we can take Astra Taylor's 'The Automation Charade' (2018), an essay whose basic thesis is that 'the rise of the robots has been greatly exaggerated' (like many authors, Taylor sees AI and robotics as pretty much synonymous, an unfortunate gloss we discuss later). Taylor agrees with our point that technological change, and automation in particular, is not a neutral process, but rather wielded from a position of class power. However, her main argument is that it is not just the actuality of automation, but more its possibility, that is weaponized to intimidate workers. She cites the threats of robotized

burger-flippers and touch-screen self-service kiosks by fast food corporations trying quell the Fight for 15 minimum wage movement. Some of those threats proved hollow, and those that have been realized, Taylor points out, still leave lots of workers toiling in McDonald's. In light of this, she proposes 'making our idea of automation itself obsolescent. A new term, "fauxtomation", seems far more fitting' (2018). Socialist feminists, she suggests, have, through their close engagement with domestic toil as unwaged work, a special insight into capitalism's ineradicable dependence on human labour, even where that labour is unacknowledged, unrewarded and conducted by women and racialized minorities. She goes on to stress the way the introduction of machines has intensified, rather than eliminated, work, emphasizing the behind-the-scenes dependence of Silicon Valley's digital platforms on the invisible work of figures such as content moderators. Against this background, Taylor takes as a clarion-call moment of insight a response she reports from the famous Marxist feminist theorist, Silvia Federici, to a conference question about capital's tendency to generate 'surplus populations': 'Don't let them make you think that you are disposable.'

Many of Taylor's points are excellent; we expand on some of them later, especially in Chapter 2, where we discuss the labour conditions of AI automation. We concur that 'automation has an ideological function as well as a technological dimension', but we disagree with her overall emphasis. While the aggregate employment effects of AI and robotics are uncertain and hotly debated, dismissal of automation as a 'charade' is deeply ahistorical. Generations of workers, from hand-loom weavers to assembly line auto-workers and cold metal print-setters would testify that there is nothing 'faux' about capital's tendency to replace humans with machines. The millions of people migrating from planetary zones bypassed by analog and digital supply chains and automated factories testify to the reality of surplus populations. While Federici may have been quite rightly suggesting that we should rethink who or what should be considered socially disposable, there is no doubt that capital always *has* made people and indeed entire populations 'disposable' (which, of course, is why it has to be resisted). Shrinking from that reality at the moment when a new instalment of corporate machinic power raises such disposability to a new level, and writing it all off as bluff and hype, may be reassuring, but it is unwise, sentimental and dangerously complacent. Probably recognizing this, at the conclusion of her essay, Taylor abruptly changes course, and concedes: 'There is no denying that technologi-

cal possibilities that could hardly be imagined a generation ago now exist, and that artificial intelligence and advances in machine learning and vision put a whole new range of jobs at risk. Entire industries have already been automated into nonexistence.' And she rightly remarks that the 'emphasis on technological factors alone, as though "disruptive innovation" comes from nowhere or is as natural as a cool breeze, casts an air of blameless inevitability over something that has deep roots in class conflict' (2018). To which we say *d'accord*. But confronting these issues demands understanding AI and its automating capacities, accompanied though they are with abundant mystification and fetishization, as something more than just a 'charade'.

Our third object of critique, the left 'maximalist' position, is the diametric opposite of the 'minimalist' approach. Not only does it hold that AI and associated technologies, such as robotics, are 'for real', and have the capacity to drastically transform the conditions of production and work, it also sees these capacities as stepping stones to socialism. Proponents of this view look optimistically at the automating capacities of AI as an opportunity to ameliorate, perhaps eventually abolish, the exploitation of wage-labour, opening up prospects for a society in which people enjoy more free time, for pleasure, personal development and/or political engagement. This seems to offer a path for socialists that is more achievable than the daunting prospect of a full-scale revolution against capital. Instead, it can be attained by a social democratic government prepared both to foster the technologies of the fourth industrial revolution and to introduce a 'universal basic income' (UBI) or 'citizens' income' – a guaranteed payment to all citizens independent of any waged job. Lenin famously wrote that communism equals 'Soviets plus electrification'. It is fair to say that 'AI plus UBI' has become the formula for techno-progressive social democratic thought. A constellation of thinkers has formed around this attractor, articulated in works such as Nick Srnicek and Alex Williams's *Inventing the Future: Postcapitalism and a World Without Work* (2015); Paul Mason's *Postcapitalism* (2015) and Aaron Bastani's (2014, 2019) arguments for 'fully automated luxury communism'; the xenofeminist (Hester 2018) line of post-gender futurism; and a cluster of autonomist or *post-operaismo* theorists.

Again, we sympathize with and in many respects share the aspirations of this group; indeed, one of us has written about digital technologies in a very similar vein (Dyer-Witheford 2014), while another has suggested that Marxism might usefully incorporate similarly maximalist elements

of transhumanist thought (Steinhoff 2014). However, we have written this book in part to directly challenge some premises of such AI-optimism. In particular, we want to contest the idea that AI can easily be detached, disentangled and re-appropriated from capitalism. Here it is useful to think about the sources on which the maximalist position draws – in part from Marx's own sometimes enthusiastic embrace of the modernizing powers of the forces of production to catalyse the emergence of socialism or communism, and also the uptake and reinterpretation of this position by poststructuralist theorists such as Gilles Deleuze and Félix Guattari. Perhaps the most important source, however, is the 'accelerationist' thinking of the anti-Marxist philosopher Nick Land (2012), who uncompromisingly argues for and celebrates what he sees as the unstoppable and species transforming (or terminating) power of computation. Indeed, the inaugurating document for the maximalist line of thought we have mapped here is Williams and Srnicek's 'Accelerationist Manifesto' (2013), which attempted a leftist re-do of Land's thought: the most accurate shorthand for the group of 'maximalist' theorists we have described is 'left accelerationists'.

As we will argue at length later, this appropriation of Landian thought dodges some of its originator's key arguments. For one, Land (2014) held AI to be the consummatory technology of capitalism, one that implanted the logic of capital at its very core. AI, in Land's view, is not merely appropriated by capital, but constituted by it: it is a technology made from and for its processes of labour automation, commodity acceleration and financial speculation. A second, yet more disquieting Landian point, is that this mutual embedment of capital and AI leads not to human emancipation from capitalism, but, on the contrary, to capital's emancipation from the human: a capital that no longer needs *homo sapiens*; human extinction.[1] These are not comfortable thoughts. And they are made even less comfortable by the fact that Land in his recent writings has emerged as a reactionary champion of racist and misogynist 'dark enlightenment' ideas that have in complex ways infiltrated the culture of Silicon Valley where much AI production takes place.[2] While we emphatically disassociate ourselves from this aspect of Landian thought, we nonetheless believe that any communist position on AI has to take his original accelerationist proposition – that AI has an elective affinity with capitalism and is fundamentally inhuman – seriously.

Given our anti-capitalist critique of both left minimalist denial and maximalist celebration of AI, it might be expected we go on to enunciate

some middle-ground, moderate position. This is not the case, or true only in the sense that we want to remove the floor beneath both minimalist complacency and maximalist optimism. Our critique of AI can best be characterized as 'abyssal', and this in two senses. First, we confess, as we think other AI thinkers should, that there are vast indeterminacies about the directions and destinations of AI-infused capitalism. Peering into the conflicting estimates of AI's near and far future capacities and deployments can, and should, instil political vertigo. The everyday uses of AI now commonplace in advanced capitalism give some indicators of its future trajectory, but no certainties. This may seem an odd assertion for a Marxist theorization of AI, given that Marxism has in many incarnations asserted bold teleological certainties; however, as we argue later, Marx's work itself contains divergent accounts of the outcome of capital's technological compulsions. We read it as a matrix of possibilities, rather than a promissory note. In and of itself, this approach undercuts complacencies both that social struggles persist unchanged, regardless of new technologies, or, conversely that, because of the same new technologies, capital's self-destruction is imminent.

That said, however, the second 'abyssal' aspect of our AI analysis is that amongst the maze of future possibilities, some potential outcomes can be discerned that are far more deeply disturbing than is allowed by either the maximalist or minimalist positions, with their respective confidence about the continuation or the end of capitalism. These outcomes throw into question assumptions about the labour theory of value, the continued centrality of struggles at the point of production, or even the confidence that capitalism cannot survive the abolition of its human waged workforce. These points demand consideration, not to justify defeatism, but as a component of a revival of revolutionary communist thought. This is what we mean when we say that, at the same time as making a Marxist critique of AI, we make an AI-informed critique of Marxism. What then is AI?

'A MACHINE CAN BE MADE TO SIMULATE IT'

To understand the effervescence surrounding AI we need to define what AI is and how it functions. We are, emphatically, not AI experts; we will make errors deriving from our lack of technical knowledge as well as from the rapidly evolving nature of the field. Despite such difficulties, we believe grappling with basic AI concepts and how AI actually works

is important. Too many accounts of AI, celebratory or dismissive, skip this effort. But it is only through some familiarity with the science and technology of AI that an effective critique can be mounted.

The workshop at Dartmouth College in Hanover, New Hampshire in 1956 is usually taken as the start of the field and study of 'artificial intelligence'. The organizers described their goal as follows:

> The study is to proceed on the basis of the conjecture 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. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. (McCarthy et al. 1955)

Since then, definitions of AI have been many and vague. AI experts show no consensus (Faggella 2018c). Compounding this definitional problem is the 'AI Effect', whereby as soon as AI can do something, it is no longer considered to require intelligence. Pamela McCorduck noted that in the history of AI 'every time somebody figured out how to make a computer do something – play good checkers, solve simple but relatively informal problems – there was chorus of critics to say, "that's not thinking"' (2004: 204). One recent AI textbook quotes Elaine Rich's pithy definition of AI from 1983: 'the study of how to make computers do things at which, at the moment, people are better' (quoted in Ertel 2018: 2). A more formal definition of AI we find useful is:

> The essence of AI – indeed, the essence of intelligence – is the ability to make appropriate generalizations in a timely fashion based on limited data. The broader the domain of application, the quicker conclusions are drawn with minimal information, the more intelligent the behaviour. (Kaplan 2016: 5–6)

This definition distinguishes AI from mere computation and allow us to differentiate between different types of existing and hypothetical AIs by considering their speed, quantity of information required, and generality of application. Kaplan's definition, however, says nothing about what AI looks like out in the world. AI does *not* mean robot, a confusion that can be blamed on pop culture. The roboticist Alan Winfield offers three complementary definitions of a robot:

1. an artificial device that can *sense* its environment and *purposefully act* on or in that environment;
2. an *embodied* artificial intelligence; or
3. a machine that can *autonomously* carry out useful work (2012: 8)

The most important aspect of Winfield's definitions is that, despite differing morphologies, all robots have bodies. AI, however, is software and, therefore, need not be embodied, though it requires computing hardware to run on. Advanced robots employ AI for functions including perception, planning actions, and learning, but a robot body does not necessarily entail AI, nor does an AI system necessarily entail a robot body.

To distinguish actually-existing AI from its speculative future incarnations, it is helpful to employ the following three categories: narrow AI, artificial general intelligence (AGI), and artificial superintelligence (ASI). Actually-existing AI is narrow: 'the vast majority of current AI approaches … are primarily designed to address narrow tasks' (Johnson et al. 2016: 4246). Most AI research, all commercial applications of AI, and the AI that consumers use daily, are such task-based tools. They are functionally more akin to microscopes than the anthropomorphic and politically active droid L3–37 in *Solo*. These systems have none or very little ability to do anything beyond their particular domain of functionality. An AI system that recognizes faces in photographs is not going to be able to process recordings of speech, play Go, or compose emails, and it is definitely not going to be able to speak Farsi. We will discuss dozens of existing narrow AI systems over the course of this book.

On the basis of generality, narrow AI is contrasted with artificial general intelligence (AGI), which refers to AI with 'the capacity for efficient cross-domain optimization' or 'the ability to transfer learning from one domain to other domains' (Muehlhauser 2013). AGI refers to an AI with the capacity to engage and behave intelligently in a wide variety of contexts and to apply knowledge learned in one context to novel situations, meaning it would be 'capable of reasoning across many intellectual domains' (Baum 2018a: 3). As of 2019, AGI remains a speculative technology, although serious research is now being conducted on it in both public and private institutions. We discuss AGI in Chapter 3.

Artificial superintelligence (ASI) is yet more speculative. While an ASI 'is likely to have general intelligence' (Baum 2018a: 3), it specifically refers to an AI 'that greatly outperform[s] the best current human minds

across many very general cognitive domains' (Bostrom 2014: 63). ASI is a science fiction staple, but serious discussion of it also occurs in academic circles where it is often seen as swiftly following the creation of AGI (see e.g. Bostrom 2014; Torres 2018; Baum 2018a; 2018b). Most commonly, the scenario imagined is that an AGI gains the ability to self-modify and evolves into a god-like ASI with unpredictable powers. The consequences of such an event are impossible to predict with certainty, but the mere possibility of it occurring compels thinkers and institutions – including Nick Bostrom and the Future of Humanity Institute at Oxford, Seth D. Baum and the Global Catastrophic Risk Institute, as well as Eliezer Yudkowsky and the Machine Intelligence Research Institute (MIRI) – to argue that we must seriously research the possibility now.

Another important distinction is that between 'strong' and 'weak' AI. While sometimes the term strong AI is used to refer to AGI (Kurzweil 2005a: 260), the term originally derives from the work of the philosopher John Searle (1980), who used it to describe the position of those who believe that an advanced AI would be conscious. Searle critiques this view from the sceptical position of weak AI, which holds that machines can never be conscious. Searle's famous Chinese Room thought experiment, which hypothesizes a human (or machine) equipped with an encyclopaedic set of rules for translating Chinese into English without being able to speak or understand the language, attempted to prove this position. As Kaplan puts it, 'strong AI posits that machines do or ultimately will have minds, while weak AI asserts that they merely simulate, rather than duplicate, real intelligence' (2016: 68). We do not take a definite stance on the question of machine consciousness. The arguments in this book do not depend on machine consciousness being physically or even logically possible, nor on the impossibility of such. On this topic we are functionally agnostic, and, as we argue in Chapter 3, so too is capital.

Actually-existing narrow AI is typically divided into three schools of thought: Good Ol' Fashioned AI (GOFAI), machine learning (ML), and the situated, embodied and dynamical framework (SED). GOFAI, also known as symbolic AI, was the first approach to AI and remained dominant until the 1980s (Boden 2014: 89). It is an approach that aims to implement high-level cognitive functions, such as logical reasoning, in machines through the manipulation of information encoded in a symbolic language. Such a system creates internal representations of its world or a problem domain in a symbolic language and performs logical

manipulations on this representation to think or act. These systems are often constructed out of sets of clearly defined rules. The best examples of GOFAI are so-called 'expert systems' or 'knowledge systems', which emerged and proliferated in the 1980s. These were intended to capture the knowledge of human experts and make it available to less skilled workers or ignorant managers. Expert systems were used for medical diagnosis, credit scoring and analysis, and business management, but the most famous is IBM's chess-playing system Deep Blue, which in 1997 defeated the reigning world champion Garry Kasparov. However, GOFAI required vast sets of rules with myriad possible interactions. Solving complex problems in this way necessitates tremendous computational power; for these and other reasons, approaches other than GOFAI were pursued.

One reaction to the problems of GOFAI is the 'situated, embodied, dynamical (SED) framework' (Beer 2014: 128); Rodney Brooks, a pioneer in the field, called his approach 'nouvelle AI' to emphasize its qualitative break with GOFAI (Copeland 2000). SED researchers are often motivated by 'Moravec's paradox' – the observation by roboticist Hans Moravec that 'it is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility' (1988: 15). SED approaches to AI emphasize the irreducible importance of the body – with its perceptual apparatuses and morphology – to cognition: for this school, it is through solving material problems that machines can evolve intelligent behaviours. Such approaches are therefore often concerned with robotics and artificial life in addition to AI. SED practitioners initially focused on very simple, insect-like robots, but in the 2010s more complex, partially humanoid robots became possible and have been introduced into industrial settings. It is possible that some variety of the SED framework could be the next dominant AI paradigm.

Another reaction to GOFAI was machine learning (ML). The ML school, formerly called connectionism, existed as early as the Dartmouth workshop, gained some traction in the 1980s with advances in learning algorithms, but did not explode until the 2010s when big data and cheap computing power proliferated. As of 2019, ML is the dominant approach to AI.[3] It is a statistical pattern-recognition approach. One NVidia researcher has described ML as a process comprised of three steps: '(1) take some data, (2) train a model on that data, and (3) use the trained