CHAPTER 8

The Social Reconfiguration of Artificial Intelligence: Utility and Feasibility

James Steinhoff

Introduction

This chapter addresses the notion of 'AI for everyone' via the concept of social reconfiguration. This was originally formulated by Jasper Bernes (2013) in his critique of what he calls the 'reconfiguration thesis' or the assumption, held by many Marxists and other critics of capital, that 'all existing means of production must have some use beyond capital, and that all technological innovation must have … a progressive dimension which is recuperable'. In other words, existing technologies which have been produced by capital for the advancement of capitalist industry can and should be appropriated and redirected towards non-capitalist, democratically-determined and socially-beneficial ends – the means of production can and should be seized.

The reconfiguration of AI is a timely topic because, since 2015, almost all the big USA tech companies, such as Google and Microsoft, have announced commitments to the so-called 'democratisation' of AI. Critiques of such programs have already been provided (Garvey 2018; Sudmann 2020; Dyer-Witheford, Kjøsen and Steinhoff 2019, 52–56) and Marxists have long pointed out that capitalism is defined by technology not being democratically controlled, but rather designed and deployed to serve the interests of one small sub-group of the world population, the owners of capital (Braverman 1998).

How to cite this book chapter:
Steinhoff, J. 2021. The Social Reconfiguration of Artificial Intelligence: Utility and Feasibility. In: Verdegem, P. (ed.) *AI for Everyone? Critical Perspectives.* Pp. 123–143. London: University of Westminster Press. DOI: https://doi.org/10.16997/book55.h.

In this chapter, I schematise the concept of reconfiguration with two dimensions: utility and feasibility. I argue that existing considerations of the reconfiguration of AI have primarily focused on utility and have largely neglected questions of feasibility. Feasibility is considered primarily in relation to the materiality of AI, or its concrete aspects which 'set constraints on and offer affordances for use' (Leonardi and Barley 2008, 171). By attending to the materiality of AI we can see how it differs from traditional, industrial means of production.

The chapter first discusses the contemporary form of AI called machine learning and its increasing importance to the tech industry. Then I discuss several aspects of its materiality. Next, I discuss Marxist theories of technology and existing evaluations of reconfiguring AI, which focus primarily on utility. Then I turn to the question of feasibility. I conclude that the social reconfiguration of AI faces substantial difficulties posed by the lack of visibility and non-modularity of AI, but that some promise is to be derived from the data commons movement. I suggest that further research on socially reconfiguring technology should focus more on feasibility, rather than utility and can begin by looking at concrete ways to resist the impositions of AI capital.

Machine Learning Materiality

Industry

Early approaches to AI attempted to automate high-level logical reasoning represented in formal languages. Such approaches to AI are called 'symbolic' or 'good old-fashioned' AI (Haugeland 1989) and have largely been overshadowed by a different approach to AI known as machine learning. Machine learning is often anthropomorphised, but it is at base the use of statistical methods, called learning algorithms, to find patterns in large datasets. On the basis of these patterns an algorithm called a 'model' is produced which may be used to analyse new data (Alpaydin 2014, 2–3). A model thus represents 'knowledge' of the patterns found by the learning algorithm and can be used to make useful analyses or predictions. Much of the hype around machine learning derives from this automated production of models from data, which Domingos (2015) calls the 'inverse of programming' (6–7).

Machine learning is being applied almost anywhere electronic data is accessible. Brynjolfsson and McAfee (2017) argue that machine learning is a general-purpose technology comparable to the combustion engine. While this remains to be seen, AI has found diverse applications from recommendation engines and targeted advertising to predictive policing software, predictive maintenance, customer resource management and fraud detection. Capital became visibly interested in machine learning around 2015. All the biggest tech companies in the world have since shifted to AI-intensive directions, including

Google, Amazon, Microsoft, Facebook, IBM, Baidu, Alibaba and Tencent. Older industrial capitals like Siemens and General Electric have followed suit. In addition to these huge companies are a variety of middle sized companies and an array of startups. Investment in AI startups increased from $1.3 billion in 2010 to over $40.4 billion in 2018 (Perrault et al. 2019, 6). This money trickles down to some, but not all, workers involved in producing AI. Salaries for machine learning scientists and engineers average $100,000 to $150,000 USD, with lavish benefits (Stanford 2019) while essential data-preparing 'ghost workers' are precariously employed through platforms like Amazon Mechanical Turk and are minimally remunerated (Gray and Suri 2019; Li 2017).

Data

Perhaps the most fundamental aspect of the materiality of machine learning is that it requires a lot of data from which to extract patterns (Alpaydin 2014, 1–4). One can get an idea of the requisite quantities by looking at some popular datasets. The dataset MNIST, a collection of handwritten digits, contains 70,000 images. Compare it to ImageNet, comprising 14,197,122 images labelled with categories and subcategories. The category 'person' has 952,000 images and 2,035 subcategories (ImageNet 2010). ImageNet is dwarfed by the Gmail Smart Reply training set which contains 238,000,000 examples, and the Google Books Ngram set which amounts to 468,000,000,000 examples. Google Translate is said to employ a dataset numbering somewhere in the trillions (Google 2019). Machine learning is no more than the sophisticated recognition of patterns across such large datasets (for a sober walkthrough of this process see Broussard (2018, 87–120)).

The companies that produce machine learning commodities are unsurprisingly concerned with obtaining vast quantities of diverse data. It is no coincidence that the major producers of AI operate a variety of platform business models in which, by acting as intermediaries between users, they can appropriate all kinds of data (Srnicek 2017). However, quality, as well as quantity, of data is important. Data does not come ready-to-use and requires labour intensive formatting, cleaning and labelling (Gitelman 2013; Gray and Suri 2019).

Compute

Producing machine learning systems requires powerful computing hardware. Since few companies can afford to buy such hardware, most advanced machine learning models are trained and deployed through the cloud platforms of the tech giants. Amazon Web Services dominates the market, but Google, Microsoft, IBM and Baidu all have their own cloud platforms. Computational power required for both training and deploying machine learning models is continually increasing. The amount of computing power used in the largest AI training runs increased 300,000 times from 2012 and 2018, with no end in sight (Amodei and Hernandez 2018).

Such computation is energy intensive. Cloud platforms thus rely on access to energy infrastructures. According to Pearce (2018), the largest data centres consume as much power as a city of a million people, and in total, data centres consume 'more than 2 percent of the world's electricity and emit roughly as much CO_2 as the airline industry'. Google, Microsoft and Baidu derive 15%, 31% and 67% of their energy, respectively, from coal (Cook 2017, 8). Efforts to 'green' the cloud by increasing renewable energy sources are ongoing, but many such campaigns consist of offsetting or buying carbon credits and do not actually mean that clouds are contributing less to CO_2 production.

Distribution

Machine learning requires sources of data, such as social media platforms. Neither can it function without storage for data, the substantial work which goes into preparing data, nor the cloud or energy infrastructures. Contemporary AI is thus not a discrete technological artifact. Even a relatively simple AI product, such as a smart home speaker, draws on a 'vast planetary network' of labour, data, resources and technologies (Crawford and Joler 2018). Machine learning cannot be analytically separated from the globally distributed infrastructure, both technical and human, on which it relies. And it is perhaps on its way to itself becoming another layer of infrastructure. Science and technology studies scholars have demonstrated how as infrastructures mature, they become 'ubiquitous, accessible, reliable, and transparent' (Edwards et al. 2007, i). Although machine learning as yet possesses none of these qualities perfectly, it is integrated into many people's daily lives in ways that increasingly approach them. Advocates of the AI industry are already positioning AI as a utility comparable to electricity or the internet, a kind of 'cognition on tap' (Steinhoff 2019). This effort to cast AI as something immediately available everywhere for users is complemented by a technological effort, dubbed 'democratisation', to make the production of AI available to a wider range of users.

'Democratisation'

The tech giants have, since around 2015, extolled the 'democratization' of AI. According to Microsoft CTO Kevin Scott, this means 'making sure that everyone has access to those platforms so that they can use the techniques of AI to enhance their own creativity, to build their own businesses, to do their jobs' (Agarwal 2019). For Madhusudan Shekar, Head of Digital Innovation at Amazon Internet Services, the democratisation of AI 'is about making the tooling and capabilities of AI/ML available to developers and data scientists at various levels of competence so that anybody can use AI to increase the velocity of new customer acquisition, reduce costs and look for new business innovations' (quoted in Ananthraj 2019). In general, democratisation efforts take the form of more or less automated, cloud-based tools and libraries (some of which are free or open source) which either assist developers with building machine learning models or help unskilled users incorporate premade models into other media. A much-feted early episode of democratisation occurred in 2015 when Google open-sourced its now widely-used TensorFlow library.

Garvey (2018) correctly points out that claims of a 'democratization' of AI draw on 'explicitly political' language but do not specify whether or how such programs are 'informed by political conceptions of democratic governance' (8079). The idea appears to be that simply distributing the tools constitutes the democratisation of AI. But this neglects consideration of how the AI products which people encounter in their daily lives are not produced or deployed through processes of democratic deliberation. Nor do such formulations address the larger issue of how the capitalist mode of production itself is premised on the exclusion of certain stakeholders from social decision-making processes. Marxists have discussed this via the distinction between the capitalist (who own and control the means of production) and working classes (who own only their ability to labour). A real democratisation of AI would require that not only capital controls its development and deployment.

Capital, Labour and Machines

Marx held that the relation between the capitalist and working classes was necessarily antagonistic. Capital has one primary goal: to increase. Marx (1990) called this valorisation (293). Mechanisms for valorisation vary across the historical permutations of capitalism, but all rely on the exploitation of labour and the capture of surplus-value. While capitalists and functionaries of capital may argue that the valorisation of capital is co-extensive with social flourishing, the COVID-19 pandemic of 2020 laid the antagonism bare, with CEOs like Jeff Bezos and political functionaries like former US President Trump openly willing to sacrifice workers for the generation of surplus-value. It appears increasingly obvious that, as Land (2017) has argued, capital has 'no conceivable meaning beside self-amplification'.

While the interests of labour in the context of work typically take the form of better wages and working conditions, the broader interests of labour are the interests of socially-existing humans as such and are therefore not amenable to *a priori* description. Marx (1993) describes the ultimate interests of labour as the 'absolute working-out' of 'creative potentialities' (488). Class antagonism results from labour's broad horizon of self-development encountering the narrow logic of capital. Labour might flourish in any number of ways not conducive to valorisation, so capital 'systematically selects for human ends compatible with its end … and systematically represses all human ends that are not' (Smith 2009, 123). One of its most effective means for doing so are machines. Driven by competition and class struggle, capital introduces machines to increase productivity by cheapening labour power, increasing control over, and dispensing with, labourers.

But this is not to say that Marx considered technology inherently opposed to labour. On the contrary, Marx held that human flourishing could only be achieved with the help of machines. For Marx and Engels (1969), communism only becomes possible when 'machinery and other inventions [make] it possible to hold out the prospect of an all-sided development, a happy existence, for all members of society'. However, before communism can be attempted, machinery must be wrested from capital. Workers must seize the means of production, or 'overthrow … the capitalists and the bureaucrats [and] proceed immediately … to replace them in the control over production and distribution, in the work of keeping account of labor and products' (Lenin 1918). Machines thus are neither inherently wedded to capital nor labour, but are rather a medium for their antagonistic relation.

Since the valorisation of capital can and often does run orthogonal to the interests and wellbeing of labour and since most AI research and production today is conducted by capital, one can assume that it is largely conducted in accord with the exigencies of valorisation. In other words, AI predominantly takes a commodity form (i.e., is designed as something which can be sold for profit) or a form which can otherwise augment the valorisation of capital (i.e., harvesting user data for inputs). There is no reason to assume that AI as a means for capital valorisation stands to benefit society beyond capital. Therefore, consideration of 'AI for everyone' needs to consider how control over AI might be taken away from capital and transferred to a democratic public. If AI is to be directed towards democratically determined ends, it will first have to be seized, in the sense that Marxists have talked of seizing the means of production.

Reconfiguration and Artificial Intelligence

Bernes (2013) defines the 'reconfiguration thesis' as the assumption that 'all existing means of production must have some use beyond capital, and that all technological innovation must have … a progressive dimension which is recuperable'. Bernes first raised the notion of reconfiguration in an analysis of capital's logistics networks. In the course of his argument, Bernes interweaves the increasingly logistical nature of capitalism, critical theory, and how it can arise from workers who inhabit logistical sites of struggle. In stark contrast to his wide-ranging discourse, I will focus narrowly on the notion of reconfiguration. We can schematise reconfiguration with two dimensions: utility and feasibility.

Utility

A first step to thinking about the potential utility of a reconfigured technology is to consider how it is useful now, and to whom. Bernes (2013) argues that logistics is 'capital's own project of cognitive mapping' because it allows capital to keep track of its dispersed moving parts. It enables a new emphasis on circulation characterised by practices such as outsourcing, just-in-time production and the global arbitrage of commodities, including labour-power. It allows the segmentation and stratification of labour, and the brutal creation of 'sacrifice zones' free of labour regulations (Hedges and Sacco 2014). The utility of logistics for capital is thus 'exploitation in its rawest form' (Bernes 2013). This is not likely a use-value for a socially reconfigured AI.

Andrejevic (2020) argues that under capital, what he calls 'automated media' (including AI) tend towards the automation of subjectivity itself (129). Andrejevic argues that this is ultimately impossible on psychoanalytic grounds, but the argument that the ultimate end of capitalist AI is the emulation of subjectivity has been advanced by others. Land (2014) holds that capital and artificial intelligence possess a 'teleological identity' and that a perfected capitalism will dispense with human labour for a full-machine economy. Such speculations range afield from this paper, but they reinforce the more immediate utility of AI for capital. AI is an automation technology with diverse applications for reducing and/or eliminating labour costs and implementing new forms of control over labour processes and social relations. It was these use-values for capital that the earliest Marxist analyses of AI reacted to. In the 1980s, AI was first commercialised in the form of 'expert systems' intended to capture and automate the knowledge and reasoning of skilled workers (Feigenbaum, McCorduck and Nii 1989). Most Marxists of this era were not interested in reconfiguring AI. The near consensus was that AI heralded a new wave of deskilling and concomitant automation, aimed at cognitive, as well as manual, forms of labour (Cooley 1981; Athanasiou 1985; Ramtin 1991).

Planning

However, another strand of Marxist thought saw utility in reconfigured technologies of automation like AI and cybernetics. Both the USSR (Peters 2016) and socialist Chile (Medina 2011) attempted to apply cybernetics to solve the 'socialist calculation problem', as the economist Ludwig von Mises described it. Von Mises (1935) contended that the distribution of resources in a planned economy requires an infeasible amount of calculation and that a capitalist market economy achieves this automatically through the market and price system. While the attempts at planned economies by Chile and the USSR failed due to the primitive computers available at the time, some Marxists have continued to pursue the idea of automated economic planning.

Cockshott (1988) argued that heuristic processing techniques 'developed in artificial intelligence can be applied to solve planning problems with economically acceptable computational costs' (1). More recently he has described big data and supercomputers as the 'foundations of Cyber Communism' (Cockshott 2017). Others have pointed out that algorithmic technologies for processing vast quantities of economic data have already been developed by large corporations like Walmart and Amazon (Jameson 2009; Phillips and Rozworski 2019). Beyond the processing of economic data, Dyer-Witheford (2013) has suggested that AI could be used to lessen bureaucratic burdens: democratic processes might be 'partially delegated to a series of communist software agents … running at the pace of high-speed trading algorithms, scuttling through data rich networks, making recommendations to human participants … communicating and cooperating with each other at a variety of levels' (13).

Bernes (2013) argues that such positions assume that 'high-volume and hyper-global distribution' possess 'usefulness … beyond production for profit'. For instance, a society not structured around commodity production would not be driven to implement planned obsolescence, so one can imagine that the overall volume of things that need to be shipped across the world would decrease substantially. In addition, more localised systems of production might obviate much of the need for vast planning techniques. The broader point is that the utility of a given existing technology for socially-determined, non-capitalist ends is not a given if it was built by capitalist firms to advance valorisation. Utility therefore 'needs to be argued for, not assumed as a matter of course' (Bernes 2013).

Full Automation

Some Marxists have also speculated on the use of AI to eliminate work. This line of thought derives from Marx's notion that 'the true realm of freedom' has its 'basic prerequisite' the 'reduction of the working-day' (Marx 1991, 959). Thinkers in the USSR held that automation had a 'crucial role in the creation of the material and technical basis of communist society' (Cooper 1977,152). Since the mid-2010s, a group of Marx-influenced thinkers referred to variously as left accelerationism (Srnicek and Williams 2015), postcapitalism theory (Mason 2016) and fully automated luxury communism (Bastani 2019) have renewed support for such ideals. I refer only to the left accelerationists here, but all of these thinkers are united in calling for full automation.

Left accelerationists argue that under capital, 'the productive forces of technology' are constrained and directed 'towards needlessly narrow ends' (Williams and Srnicek 2014, 355). The technology developed by capital should be seized: 'existing infrastructure is not a capitalist stage to be smashed, but a springboard to launch towards post-capitalism' (Williams and Srnicek 2014, 355). They hold that 'existing technology [can be] repurposed immediately' (Srnicek and

Williams 2015). Alongside decarbonising the economy, developing renewable energy sources, cheap medicine and space travel, they advocate 'building artificial intelligence' (Srnicek and Williams 2015). For left accelerationists, a reconfigured AI is useful primarily in that it could contribute to full automation, which is desirable because 'machines can increasingly produce all necessary goods and services, while also releasing humanity from the effort of producing them' (Srnicek and Williams 2015). Eventually, a 'fully automated economy' could 'liberate humanity from the drudgery of work while simultaneously producing increasing amounts of wealth' (Srnicek and Williams 2015, 109).

The automation of bad work and the administration of a planned economy are certainly useful applications of AI that extend beyond the logic of valorisation. But utility should be considered alongside feasibility.

Feasibility

Even if a given capitalist technology presents useful possibilities, it is not necessarily the case that its social reconfiguration appears feasible. Bernes presents several reasons why a social reconfiguration of logistics is infeasible, two of which derive from its distributed nature, and are likewise applicable to contemporary machine learning. The first of these pertains to visibility.

Visibility

Logistics comprises a vast, heterogeneous network of technologies and institutions which remains invisible as a whole to the workers who populate its variegated zones. The means of logistical production are distributed across this network, but '[o]ne cannot imagine seizing that which one cannot visualise, and inside of which one's place remains uncertain' (Bernes 2013). Logistics is capital's means for knowing itself, but this knowledge is barred from workers. This sense of visibility is not only an issue when considering the initial seizure of a technology, but also for tracking the progress of its social reconfiguration, which is unlikely to occur instantly. To persevere, 'struggles need to recognise themselves in the effects they create, they need to be able to map out those effects … within a political sequence that has both past and future, that opens onto a horizon of possibilities' (Bernes 2013). Contemporary AI presents similar problems of visibility, from several different angles.

AI is temporally and physically distributed across layers of infrastructures. To 'see' AI we need ways to chart this vast network and make it appear as a coherent collection of people and things. Excellent work on visualising AI has been done in visual essays by Crawford and Joler (2018), which reveals the diverse materiality of AI, and by Pasquinelli and Joler (2020) which aims to 'secularize' AI by casting it not as alien intelligence, but something more like an optical instrument, akin to a microscope. But visual essays can only go so far. In a more fundamental sense, visibility is a problem of knowledge.

'Democratisation' of AI programs aim to make AI accessible to less skilled users, but they do so by abstracting from the underlying code with user-friendly interfaces. Of course, all computing technology today uses layers of abstraction, whether to allow skilled users to achieve complex ends more easily or to allow less-skilled users to do something at all (including the word processor I am using to write this chapter). Not many people write machine code. But as Kittler (1995) pointed out, increasing layers of abstraction from the underlying materiality of the computer mean that the potential ends it might be put to are reduced; layers of software act as a 'secrecy system' blocking access to basic functionality. So-called 'democratised' machine learning does not enable the production of novel applications of the technology beyond pre-determined bounds. At best, it allows more users to apply pre-canned software tools.

Further, while the open sourcing of AI tools and libraries like Google's TensorFlow may seem like a truly democratic move insofar as companies are giving away proprietary software, it also has competitive dimensions motivated by valorisation. Open sourcing can generate a community around the software which entails skilled developers (and potential future employees) for the company who produces the software. It can also create a software ecosystem based on those tools, which a company can retain control over through a variety of mechanisms from mandatory lock-in agreements to closed source variants of programs. Google used (and uses) such strategies to make Android the most popular mobile operating system in the world (Amadeo 2018). While Google's TensorFlow can currently be run on competing clouds, there are indications that the tech giants are aiming towards fully siloed AI ecosystems. Google is not alone in developing proprietary hardware specially designed for AI. Google's Tensor Processing Unit (TPU) provides a 'performance boost' over traditional hardware, but 'only if you use the right kind of machine-learning framework with it … Google's own TensorFlow' (Yegulalp 2017). Open source AI software is thus one tactic of a larger strategy by which AI capitals combat their rivals for a share of surplus-value.

Visibility is also a technical problem. Machine learning has a 'black box' problem because the complex computations that occur within a system cannot be disassembled and examined and thus its outputs remain inexplicable. As one researcher puts it, the: 'problem is that the knowledge gets baked into the network, rather than into us' (quoted in Castelvecchi 2016). Even if some machine learning models could be reconfigured without being rebuilt, their operations would remain inscrutable, presenting problems of accountability (Garigliano and Mich 2019). The delegation of economic planning or bureaucratic decision-making to a black box might be tolerable for some, as long as no mistakes are made, but it seems dubious that such occult mechanisms would represent a substantial improvement for democratic decision-making over delegating social decisions to the so-called logic of the market.

Non-modularity

A second dimension of feasibility also concerns distribution, but from a tactile, rather than visual, standpoint. Bernes (2013) argues that while revolutions are necessarily localised, 'any attempt to seize the means of [logistical] production would require an immediately global seizure'. Without connection to the rest of the logistical network, a reconfigured port facility is of little use. On the other hand, maintaining connection with the rest of the network entails 'trade with capitalist partners, an enchainment to production for profit … the results of which will be nothing less than disastrous' (Bernes 2013). One might reply that taking the whole system over at once is not necessary – one can appropriate it piecemeal. This might be the case, but it needs to be taken into account that infrastructures are built on top of infrastructures and intertwined with them in 'recursive' ways (Larkin 2013, 30). A technology that is part of a larger system may not necessarily be possible to reconfigure by itself.

In a second consideration of the problem of reconfiguration, focused this time on agriculture and energy, Bernes discusses the non-modularity of certain technologies. By this he means technologies that 'fit together into technical ensembles that exhibit a strong degree of path-dependency, meaning historical implementation strongly influences future development, precluding or making difficult many configurations we may find desirable' (Bernes 2018, 334). He singles out energy infrastructure as particularly non-modular and argues that hopes of simply substituting clean energy sources, even if all political opposition were removed, is wishful thinking because the 'technology [we] would inherit works with and only with fossil fuels' (Bernes 2018, 334).

To consider the non-modularity of machine learning, recall its reliance on the highly centralised clouds maintained by the tech giants. Any reconfiguration of AI would require a seizure of the data centres which make up the cloud as well as the energy sources and infrastructures necessary to power them. Such facilities could, certainly, be seized like more traditional means of production, such as factories. But this presents its own host of material problems. One concerns the powerful hardware required for AI and its energy consumption. While some greening of data centres is evidently possible, it is uncertain whether greening efforts can keep pace with the increasing computational demands of machine learning. Developers at OpenAI recently stated that 'it's difficult to be confident that the recent trend of rapid increase in compute usage will stop, and we see many reasons that the trend could continue' (Sastry et al. 2019).

Cutting-edge machine learning is increasingly out of reach for organisations without resources on par with Facebook or Google. OpenAI was founded as a non-profit research lab with substantial donations from the likes of Elon Musk and Peter Thiel, but in 2019, justified a switch to a 'limited profit' model, in partnership with Microsoft, because AI research 'requires a lot of capital for computational power' (Brockman 2019). If contemporary machine learning algorithms are indefinitely scalable, meaning that their performance improves as long as more data and computational power are made available, then the hardware cost of AI research and development will continue to rise. If the reconfiguration of AI is to occur in a local context, and if it wishes to remain on functional par with capitalist AI, it will have to devote considerable resources to the requisite hardware and figure out how to make their operation more ecologically feasible.

But perhaps seizing data centres is not necessary for the reconfiguration of AI. Some commentators hope to shift the computational load from the centralised cloud onto individual devices in a technique called decentralised or edge computing. While increasing amounts of edge computing seem likely as components continue to decrease in size, data centres will always offer more space and thus more total computing power. The expert consensus seems to be that with existing technology it is 'not possible to move Cloud-levels of compute onto the edge' (Bailey 2019). Another alternative to seizing the existing cloud could be to construct an alternative cloud. Such initiatives exist, such as the CommonsCloud Alliance, which aims to build a cloud based not on centralised data centres, but on computing power and storage space shared amongst users (Sylvester-Bradley 2018). This seems feasible, but unlikely to compare to the capacities of the clouds of the tech giants.

Data itself also raises several questions of feasibility. The first pertains to data collection. Many AI systems are trained on publicly available datasets in early stages of development, but usually, proprietary datasets are necessary to complete a project (Polovets 2015). The preparation and labelling of these is a labour-intensive and time-consuming process (Wu 2018). Creating a dataset also requires a venue for data collection in the first place. Companies such as Amazon and Google harvest reams of data from the interactions of users of their applications, even when they claim not to be, as smart home devices have shown (Fingas 2019). One business analysis of IBM suggests that because the company lacks a data collection venue, it will face difficulties developing its AI endeavours (Kisner, Wishnow and Ivannikov 2017, 19–20). This perhaps indicates why, in 2020, IBM entered into partnership with data-rich enterprise software company Salesforce. AI entails a capitalism built around surveillance, enabling 'data extractivism' (Zuboff 2019). How desirable is pervasive, multimodal surveillance for a socially reconfigured AI?

Machine learning's reliance on data also necessitates a unique form of maintenance. A model which functioned well when it was deployed will no longer do so if the domain it is applied to changes such that the data it was trained on no longer accurately reflects that domain (Schmitz 2017). Imagine a hypothetical model trained to recognise traffic signs. If overnight the red octagons reading STOP were replaced with purple triangles reading HALT, the model would no longer function and would require maintenance. A social reconfiguration of AI will presumably be one component of a larger democratic restructuring of society with substantial changes to the normal routines of social life. A preview of this sort of disruption for AI has been provided by the COVID-19 pandemic, 'models trained on normal human behavior are now finding that normal has changed, and some are no longer working as they should' (Heaven 2020). When substantial shifts in human behaviour occur, models no longer map onto reality. It is reasonable to assume that models trained on data produced by life under capital may not function in a society striving to fundamentally change its basic axioms.

There is, however, at least one reason for optimism concerning data. A promising alternative to mass surveillance and siloed data ecosystems comes from the notion of data commons, in which individuals and institutions share data willingly, with controls over anonymity and a goal to make data valuable not only to tech companies, but also to its producers. The DECODE projects in Barcelona and Amsterdam have piloted aspects of a data commons successfully and are planning to scale up in the future (Bass and Old 2020). An interesting aspect of these projects is their use of other relatively new technologies, such as smart contracts (Alharby and Van Moorsel 2017), to aggregate and analyse sensitive data in ways which preserve privacy and retain user control. These projects provide a concrete demonstration of the feasibility of reconfiguring some aspects of data ecosystems. That they draw on novel smart contracts should remind us that assessments of feasibility are necessarily contextual; they are constrained by the knowledge of the assessor and the current technological milieu. As such, this chapter makes an argument which remains open to revision. Any social reconfiguration of AI will have to go beyond the assessment attempted here and search out such novelties, technological or other, as might be relevant.

Conclusion: Counter-AI

Discussing the democratization of AI, Kevin Scott, CTO at Microsoft, makes the following comparison with the industrial revolution:

> the people who benefited from [steam powered] technology were folks who had the capital to … build factories and businesses around the machines and people who had expertise to design, build and operate them. But eventually … the technology democratized. You don't get any sort of advantage now, as a capital owner, because you can build an engine. And what we … need to do … is dramatically contract that period of time where AI is so hard to do that only a handful of people can do it. (Agarwal 2019)

Scott's sense of democratisation here hinges on the mere generalisation of a technology. The notion seems to be that because, over time, knowledge of how to build steam engines diffused through the population, this technology became democratised – any 'capital owner' can build or go out and buy a steam engine. But this formulation seems blissfully unaware of the inequalities between capital owners and labour and thus it precisely misunderstands the meaning of democratisation. There are a lot of people in the world without any capital at all. Further, the mere distribution of free AI tools does not ensure democratic control over the centralised means of AI production nor upset the advantage held by the current producers of AI. This chapter has thus explored what it might mean to actually democratise AI, or rather, to socially reconfigure existing AI into an 'AI for everyone'. The central point I have hoped to make is that consideration of the utility of a socially-reconfigured AI should be complemented by consideration of feasibility, which is largely determined by the 'material character of the powers and forces' involved in the technology (Bernes 2018, 336).

Reconfiguring AI entails simultaneous reconfiguration of large chunks of the tech sector, energy infrastructure, advertising industry, data market/ecosystem, and also requires social deliberation over aspects of the material character of AI, such as its apparent need for surveillance. This assessment resonates with that of Huber (2020), who lucidly argues that a reconfiguration of the capitalist food industry is impossible via incremental piecemeal tweaking, but will instead require revolutionising the entire system. Morozov (2019) makes the same case for the 'feedback infrastructure' or the means of producing, harvesting and processing data which are so essential to AI. On the other hand, the data commons movement indicates practical ways in which the data which machine learning systems are built from can be utlised to benefit those not at the helms of big tech capitals. While the data commons movement is occurring within the circuits of capital, it shows how a social reconfiguration of AI might begin. Even if the seizure of AI seems herculean, data commons projects demonstrate a concrete modicum of feasibility.

Finally, if one cannot seize AI today, one can still resist it. It is true, however, that resistance is a wearily overused term for critics of capitalism. What does it mean, in practice, to resist capitalist machine learning? For a final time, I will draw on Bernes (2013), who suggests that we might imagine a 'logistics against logistics, a counter-logistics which employs the conceptual and technical equipment of the industry in order to identify and exploit bottlenecks … This counter-logistics might be a proletarian art of war to match capital's own *ars belli*'. While this chapter cannot adequately explore the idea, it can suggest that there could be a proletarian counter-AI built around the axis of data on which machine learning, and the capital it increasingly powers, runs.

Early forms of this might take the form of rendering data unavailable or unusable to capital. Users might engage in 'data strikes' by deleting or otherwise denying access to their data (Vincent, Hecht and Sen 2019) and they might distort or 'poison' their data by introducing inaccurate or harmful patterns into it (Vincent et al. 2021). But what about the data infrastructure more broadly? It should be possible to determine key bottlenecks in the valorisation processes of capitalist AI, at which democratic control might be one day exercised, but for now might provide at least an opening for proto-democratic intervention. However, since secrecy is a prime virtue of AI capital, it can be difficult to obtain information on its data-intensive processes. One might thus look into the technical literature on AI for concerns which might be exploited by those resisting capitalist AI, such as 'adversarial attacks' which exploit the pattern recognition properties of machine learning to render model output inaccurate (Samangouei, Kabkab and Chellappa 2018). However, technical problems need to be considered in relation to how they are implicated within valorisation processes. Thus, a fruitful direction for research is business-oriented literature on AI adoption and production. This is generated by the producers of AI commodities – Microsoft's online AI Business School and Google's array of free AI education courses are two examples – but also by a wide variety of industry promoters, from consulting firms, see Accenture's (n.d.) guide to 'AI for Business Transformation' and management-oriented books like *The Executive Guide to Artificial Intelligence* (Burgess 2018). Such sources can reveal what AI capitals are worried about and indicate potential bottlenecks amenable to outsider intervention. Once identified, bottlenecks can be analysed and the social relations which support valorisation via AI therein might be replaced with alternative social relations not amenable to the data-hungry valorisation of AI capital. Finding bottlenecks returns us to the question of visibility, without which strategy cannot be formulated. I hope this chapter will contribute to an incremental increase in visibility and, perhaps, a half step towards a strategy.

References

Accenture. n.d. AI for Business Transformation. https://www.accenture.com/us-en/services/applied-intelligence/business-transformation

Agarwal, D. 2019. Chat with Kevin Scott EVP and CTO of Microsoft. 15 July. *Venture Beat*. Last accessed 5 May 2020: https://www.youtube.com/watch?v=npXxiVaY9p0

Alharby, M. and Van Moorsel, A. 2017. Blockchain-Based Smart Contracts: A Systematic Mapping Study. *arXiv preprint arXiv:1710.06372*.

Alpaydin, E. 2014. *Introduction to Machine Learning*. Cambridge, MA: MIT Press.

Amadeo, R. 2018. Google's Iron Grip on Android: Controlling Open Source by Any Means Necessary. 21 August. *Ars Technica*. Last accessed 5 May 2020: https://arstechnica.com/gadgets/2018/07/googles-iron-grip-on-android-controlling-open-source-by-any-means-necessary

Amodei, D. and Hernandez, D. 2018. AI and Compute. 16 May. *OpenAI Blog*. Last accessed 5 May 2020: https://openai.com/blog/ai-and-compute

Ananthraj, V. 2019. Madhusudan Shekar on how Amazon is democratizing AI. 31 July. *TechCircle*. Last accessed 5 May 2020: https://www.techcircle.in/2019/07/31/madhusudan-shekar-on-how-amazon-is-democratizing-ai

Andrejevic, M. 2020. *Automated Media*. New York: Routledge.

Athanasiou, T. 1985. Artificial Intelligence: Cleverly Disguised Politics. In: T. Solomonides and L. Levidow (Eds.), *Compulsive Technology: Computers as Culture*, pp. 13–35. London: Free Association Books.

Bailey, B. 2019. Power is Limiting Machine Learning Deployments. 25 July. *Semiconductor Engineering*. Last accessed 15 May 2020: https://semiengineering.com/power-limitations-of-machine-learning

Bass, T. and Old, R. 2020. Common Knowledge: Citizen-led Data Governance for Better Cities. Decode Project. https://decodeproject.eu/publications/common-knowledge-citizen-led-data-governance-better-cities

Bastani, A. 2019. *Fully Automated Luxury Communism*. New York: Verso.

Bernes, J. 2013. Logistics, Counter Logistics and the Communist Prospect. *Endnotes* 3. Last accessed 5 May 2020: https://endnotes.org.uk/issues/3/en/jasper-bernes-logistics-counterlogistics-and-the-communist-prospect

Bernes, J. 2018. The Belly of the Revolution: Agriculture, Energy and the Future of Communism. In: B.R. Bellamy and J. Diamanti (Eds.), *Materialism and the Critique of Energy*, pp. 331–375. Chicago: MCM.

Braverman, H. 1998. *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. New York: NYU Press.

Brockman, G. 2019. Microsoft Invest in and Partners with OpenAI to Support us Building Beneficial AI. 22 July. *OpenAI Blog*. Last accessed 5 May 2020: https://openai.com/blog/microsoft

Broussard, M. 2018. Artificial Unintelligence: How Computers Misunderstand the World. Cambridge, MA: MIT Press.

Brynjolfsson, E. and McAfee, A. 2017. The Business of Artificial Intelligence. July. *Harvard Business Review*. Last accessed 10 May 2020: https://hbr.org/cover-story/2017/07/the-business-of-artificial-intelligence

Burgess, A. 2018. *The Executive Guide to Artificial Intelligence: How to Identify and Implement Applications for AI in your Organization*. London: Palgrave Macmillan.

Castelvecchi, D. 2016. Can we Open the Black Box of AI? 5 October. *Nature News & Comment*. Last accessed 15 May 2020: https://www.nature.com/news/can-we-open-the-black-box-of-ai-1.20731

Cockshott, P. 1988. Application of Artificial Intelligence Techniques to Economic Planning. University of Strathclyde. Last accessed 5 May 2020: http://www.dcs.gla.ac.uk/~wpc/reports/plan_with_AIT.pdf

Cockshott, P. 2017. Big Data and Supercomputers: Foundations of Cyber Communism. 26–28 September. Presented at The Ninth International Vanguard Scientific Conference on 100 Years of Real Socialism and the Theory of Post-Capitalist Civilization, Hanoi, Vietnam.

Cooley, M. 1981. On the Taylorisation of Intellectual Work. In L. Levidow and R. Young (Eds.), *Science, Technology and the Labour Process Volume 2*. London: CSE Books.

Cook, G. 2017. Clicking Clean: Who is Winning the Race to Build a Green Internet? *Greenpeace*. Last accessed 5 May 2020: https://storage.googleapis.com/planet4-international-stateless/2017/01/35f0ac1a-clickclean2016-hires.pdf

Cooper, J. 1977. The Scientific and Technical Revolution in Soviet Theory. In: F. Fleron (Ed.), *Technology and Communist Culture: The Socio-Cultural Impact of Technology Under Socialism*, pp. 146–179. New York: Praeger.

Crawford, K. and Joler, V. 2018. Anatomy of an AI System: The Amazon Echo as an Anatomical Map of Human Labor, Data and Planetary Resources. *AI Now Institute and Share Lab*. Last accessed 5 May 2020: https://anatomyof.ai

Domingos, P. 2015. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake our World*. New York: Basic Books.

Dyer-Witheford, N. 2013. Red Plenty Platforms. *Culture Machine* 14, 1–27.

Dyer-Witheford, N., Kjøsen, A.M. and Steinhoff, J. 2019. *Inhuman Power: Artificial Intelligence and the Future of Capitalism*. London: Pluto.

Edwards, P., Jackson S., Bowker, G. and Knobel, C. 2007. Understanding Infrastructure: Dynamics, Tensions, and Design. Report of a Workshop on 'History and Theory of Infrastructure: Lessons for New Scientific Cyberninfrastructures'. Last accessed 5 May 2020: https://deepblue.lib.umich.edu/bitstream/handle/2027.42/49353/UnderstandingInfrastructure2007.pdf

Feigenbaum, E., McCorduck, P. and Nii, H. 1989. *The Rise of the Expert Company: How Visionary Companies Are Using Artificial Intelligence to Achieve Higher Productivity and Profits*. New York: Vintage Books.

Fingas, J. 2019. Amazon and Google Ask for Non-Stop Data from Smart Home Devices. 13 February. *Engadget*. Last accessed 15 May 2020: https://www.engadget.com/2019-02-13-amazon-and-google-continuous-smart-home-data.html

Garigliano, R. and Mich, L. 2019. Looking Inside the Black Box: Core Semantics Towards Accountability of Artificial Intelligence. In: M.H. ter Beek, A. Fantechi and L. Semini (Eds.), *From Software Engineering to Formal Methods and Tools, and Back: Essays Dedicated to Sefania Gnesi on the Occasion of Her 65th Birthday*, pp. 250–266. Cham, Switzerland: Springer.

Garvey, C. 2018. A Framework for Evaluating Barriers to the Democratization of Artificial Intelligence. Presented at *The Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*. pp. 8079–8080. Last accessed 5 May 2020: https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/viewFile/17320/16477

Gitelman, L. (Ed). 2013. *Raw Data is an Oxymoron*. Cambridge, MA: MIT Press.

Google. 2019. The Size and Quality of a Data Set. *Google Developers Machine Learning Crash Course*. Last accessed 5 May 2020: https://developers.google.com/machine-learning/data-prep/construct/collect/data-size-quality

Gray, M.L. and Suri, S. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. New York: Houghton Mifflin Harcourt.

Harvey, D. 2007. *A Brief History of Neoliberalism*. Oxford: Oxford University Press.

Haugeland, J. 1989. *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press.

Heaven, Will. 2020. Our Weird Behavior During the Pandemic is Messing with AI Models. 11 May. *MIT Technology Review*. Last accessed 5 May 2020: https://www.technologyreview.com/2020/05/11/1001563/covid-pandemic-broken-ai-machine-learning-amazon-retail-fraud-humans-in-the-loop/?truid=8f7239a3e8ff7abf667fea197c1218cd&utm_source=the_download&utm_medium=email&utm_campaign=the_download.unpaid.engagement&utm_content=05-13-2020

Hedges, C. and Sacco, J. 2014. *Days of Destruction, Days of Revolt*. Toronto, Canada: Vintage Canada.

Huber, M. 2020. Socialise the Food System. 19 April. *Tribune*. Last accessed 5 May 2020: https://www.tribunemag.co.uk/2020/04/socialise-the-food-system

ImageNet. 2010. About ImageNet. *ImageNet.org*. Last accessed 5 May 2020: http://image-net.org/about-statssrnicke

Jameson, F. 2009. *Valences of the Dialectic*. London: Verso.

Kisner, J., Wishnow, D. and Ivannikov, T. 2017. IBM: Creating Shareholder Value with AI? Not So Elementary, My Dear Watson. 12 July. Jeffries Franchise Notes. Retrieved from: https://javatar.bluematrix.com/pdf/fO5xcWjc

Kittler, F. 1995. There is No Software. *ctheory*. Last accessed 5 May 2020: http://ctheory.net/ctheory_wp/there-is-no-software

Land, N. 2014. The Teleological Identity of Capitalism and Artificial Intelligence. Remarks to the Participants of the Incredible Machines 2014 Conference. 8 March. Formerly online at this address: http://incrediblemachines.info/nick-land-the-teleological-identity-of-capitalism-and-artificial-intelligence

Land, N. 2017. A Quick and Dirty Introduction to Accelerationism. 25 May. *Jacobite*. Last accessed 5 May 2020: https://jacobitemag.com/2017/05/25/a-quick-and-dirty-introduction-to-accelerationism

Larkin, B. 2013. The Politics and Poetics of Infrastructure. *Annual Review of Anthropology*, 42, 327–343.

Lenin, V.I. 1971 [1918]. The Immediate Tasks of the Soviet Government. In: *Lenin Collected Works Volume 42* (2nd English Edition). Moscow: Progress Publishers.

Leonardi, P.M. and Barley, S.R. 2008. Materiality and Change: Challenges to Building Better Theory About Technology and Organizing. *Information and Organization*, 18(3), 159–176.

Li, F. 2017. ImageNet: Where have we been? Where are we going? ACM Webinar. Last accessed 5 May 2020: https://learning.acm.org/binaries/content/assets/leaning-center/webinar-slides/2017/imagenet_2017_acm_webinar_compressed.pdf

Marx, K. 1990. *Capital Volume One*. London: Penguin Classics.

Marx, K. 1991. *Capital Volume Three*. London: Penguin Classics.

Marx, K. 1993. *Grundrisse: Foundations of the Critique of Political Economy*. London: Penguin Classics.

Marx, K. and Engels, F. 1969. Manifesto of the Communist Party. *Marx/Engels Selected Works Volume 1*. Moscow: Progress Publishers.

Mason, P. 2016. *Postcapitalism: A Guide to Our Future*. New York: Farrar, Straus and Giroux.

Medina, E. 2011. *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile*. Cambridge, MA: MIT Press.

Morozov, Evgeny. 2019. Digital Socialism? The Calculation Debate in the Age of Big Data. *New Left Review* 116. https://newleftreview.org/issues/II116/articles/evgeny-morozov-digital-socialism

Pasquinelli, M. and Joler, V. 2020. The Nooscope Manifested: Artificial Intelligence as Instrument of Knowledge Extractivism. KIM HfG Karlsruhe and Share Lab. 1 May. Retrieved from: http://nooscope.ai

Pearce, F. 2018. Energy Hogs: Can World's Huge Data Centers Be Made More Efficient? 3 April. *Yale Environment 360*. Last accessed 5 May 2020: https://e360.yale.edu/features/energy-hogs-can-huge-data-centers-be-made-more-efficient

Perrault, R., Shoham, Y., Brynjolfsson, E., Clark, J., Etchemendy, J., Grosz, B., Lyons, T., Manyika, J., Mishra, S. and Niebles, J.C. 2019. *The AI Index 2019 Annual Report*. Human-Centered AI Institute. Stanford, CA: Stanford University.

Peters, B. 2016. *How Not to Network a Nation: The Uneasy History of the Soviet Internet*. Cambridge, MA: MIT Press.

Phillips, L. and Rozworski, M. 2019. *The People's Republic of Wal-Mart: How the World's Biggest Corporations Are Laying the Foundation for Socialism*. London: Verso.

Polovets, L. 2015. The Value of Data Part 1: Using Data as Competitive Advantage. 27 February. *Coding VC*. Last accessed 5 May 2020: https://codingvc.com/the-value-of-data-part-1-using-data-as-a-competitive-advantage

Ramtin, R. 1991. *Capitalism and Automation: Revolution in Technology and Capitalist Breakdown*. London: Pluto.

Samangouei, P., Kabkab, M. and Chellappa, R. 2018. Defense-gan: Protecting Classifiers Against Adversarial Attacks Using Generative Models. *arXiv preprint* 1805.06605.

Sastry, G., Clark, J., Brockman, G. and Sutskever, I. 2019. Compute Used in Older Headline Results. *OpenAI Blog*. Last accessed 5 May 2020: https://openai.com/blog/ai-and-compute/#addendum

Schmitz, M. 2017. Why Your Models Need Maintenance. 12 May. *Towards Data Science*. Last accessed 5 May 2020: https://towardsdatascience.com/why-your-models-need-maintenance-faff545b38a2

Smith, T. 2009. The Chapters on Machinery in the 1861–63 Manuscripts. In: R. Bellofiore and R. Fineschi (Eds.), *Re-Reading Marx: New Perspectives After the Critical Edition*, pp. 112–127. Basingstoke: Palgrave Macmillan.

Srnicek, N. 2017. The Challenges of Platform Capitalism: Understanding the Logic of a New Business Model. *Juncture*, 23(4), 254–257.

Srnicek, N. and Williams, A. 2015. *Inventing the Future: Postcapitalism and a World Without Work*. London: Verso.

Stanford, S. 2019. Artificial Intelligence (AI): Salaries Heading Skyward. 20 September. *Towards AI*. Last accessed 5 May 2020: https://medium.com/towards-artificial-intelligence/artificial-intelligence-salaries-heading-skyward-e41b2a7bba7d

Steinhoff, J. 2019. Cognition on Tap. *Digital Culture & Society*, 4(2), 89–104.

Sudmann, A. (Ed.). 2020. *The Democratization of Artificial Intelligence: Net Politics in the Era of Learning Algorithms*. Berlin: Transcript-Verlag.

Sylvester-Bradley, O. 2018. The Making of the Cooperative Cloud. 3 April. *The Open Co-op*. Last accessed 5 May 2020: https://open.coop/2018/04/03/making-of-the-coop-cloud

Vincent, N., Hecht, B. and Sen, S. 2019. 'Data Strikes': Evaluating the Effectiveness of a New Form of Collective Action Against Technology Companies. In: *The World Wide Web Conference*, 1931–1943. New York: ACM.

Vincent, N., Li, H., Tilly, N., Chancellor, S. and Hecht, B. 2021. Data Leverage: A Framework for Empowering the Public in its Relationship with Technology Companies. *arXiv preprint*. 2012.09995.

von Mises, L. 1935. Calculation in the Socialist Commonwealth. In: F.A. Hayek (Ed.), *Collectivist Economic Planning*. London: Routledge.

Williams, A. and Srnicek, N. 2014. #accelerate: Manifesto for an Accelerationist Politics. In R. Mackay and A. Avanessian (Eds.), *#accelerate: The Accelerationist Reader*. Falmouth: Urbanomic.

Wu, H. 2018. China is Achieving AI Dominance by Relying on Young Blue-Collar Workers. *Motherboard*. 21 December. https://www.vice.com/en_us/article/7xyabb/china-ai-dominance-relies-on-young-data-labelers

Yegulalp, S. 2017. Google's Machine Learning Cloud Pipeline Explained. 19 May. *Infoworld*. Last accessed 5 May 2020: https://www.infoworld

.com/article/3197405/tpus-googles-machine-learning-pipeline-explained.html

Zuboff, S. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power.* London: Profile Books.