

---

Electronic Thesis and Dissertation Repository

---

11-25-2019 2:30 PM

## Critiquing the New Autonomy of Immaterial Labour: An Analysis of Work in the Artificial Intelligence Industry

James Steinhoff, *The University of Western Ontario*

Supervisor: Dyer-Withford, Nick, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Media Studies

© James Steinhoff 2019

Follow this and additional works at: <https://ir.lib.uwo.ca/etd>



Part of the [Other Film and Media Studies Commons](#)

---

### Recommended Citation

Steinhoff, James, "Critiquing the New Autonomy of Immaterial Labour: An Analysis of Work in the Artificial Intelligence Industry" (2019). *Electronic Thesis and Dissertation Repository*. 6669.  
<https://ir.lib.uwo.ca/etd/6669>

This Dissertation/Thesis is brought to you for free and open access by Scholarship@Western. It has been accepted for inclusion in Electronic Thesis and Dissertation Repository by an authorized administrator of Scholarship@Western. For more information, please contact [wlsadmin@uwo.ca](mailto:wlsadmin@uwo.ca).

Electronic Thesis and Dissertation Repository

---

11-25-2019 2:30 PM

## Critiquing the New Autonomy of Immaterial Labour: An Analysis of Work in the Artificial Intelligence Industry

James Steinhoff

Supervisor

Dyer-Witthford, Nick

*The University of Western Ontario*

Graduate Program in Media Studies

A thesis submitted in partial fulfillment of the requirements for the degree in Doctor of Philosophy

© James Steinhoff 2019

Follow this and additional works at: <https://ir.lib.uwo.ca/etd>



Part of the [Other Film and Media Studies Commons](#)

---

## Abstract

Karl Marx theorized capitalism as a relation between labour, capital and machines. For Marx, capital, the process of self-augmenting value appropriated from human labour, is inherently driven by competition to replace labour in production with machines. Marx goes as far as to describe machines as capital's "most powerful weapon" for suppressing working class revolt. Marx, however, could not have predicted the computing machines – such as artificial intelligence – which now form the basis for an increasingly cybernetic capital. Since Marx's time, many Marxist thinkers have sought to apply or update his approach to the cybernetic era. The influential *post-operaismo* school argues that fundamental revisions to Marx's approach are necessitated by the changed nature of high-tech capital wherein arises a novel "immaterial" type of labour. Immaterial labour, the argument goes, appropriates the machines of capital and achieves a new autonomy from capital, which can no longer control labour and instead, can only attempt to capture the fruits of its autonomous productive capacities.

This dissertation's goal is to assess the validity of *post-operaismo*'s claim for a new autonomy of immaterial labour from capital. It does so by conducting an analysis of work in the contemporary artificial intelligence (AI) industry. Work in the AI Industry should be, according to *post-operaismo*, immaterial labour *par excellence*. Therefore, this dissertation answers the following research question: does work in the AI Industry evince the new autonomy from capital attributed to immaterial labour by *post-operaismo*? I argue that it does not. I mount this argument with a multimodal methodology. I employ documentary analysis and qualitative interviews with workers and management in the AI Industry to produce a history, political economy analysis and labour process analysis of the AI Industry. This is followed by a theoretical analysis which assesses the claims of *post-operaismo* by the example of the AI Industry. I argue that work in the AI Industry remains under the control of capital and that, antipodally to claims of a new autonomy of labour, this industry evinces an increasing autonomy of capital. I conclude the *post-operaismo* mistakes obsolescence for autonomy.

## Keywords

Marxism, artificial intelligence, machine learning, political economy, political economy of media, labour process theory, post-operaismo, media studies, labour studies, digital labour, immaterial labour, Marxist theory, autonomist Marxism, New Reading of Marx

## Summary for Lay Audience

This dissertation argues that while Karl Marx may have analyzed capitalism a distant two hundred years ago, his insights remain relevant today. Marx argued that capitalist businesses use machines to oppress and control workers in the interests of harvesting what he called surplus-value. Some contemporary theorists, inspired by Marx, argue that while his analysis was valuable, it needs to be updated in substantial ways to remain relevant to the social and economic systems of today, which are characterized by advanced digital technologies that he could never have anticipated, such as the internet and artificial intelligence. This dissertation argues against one such school of thought, called *post-operaismo*, which holds that digital technologies mean that work increasingly takes a form they call “immaterial labour” in which workers gain increasing control over their own work, ultimately leading to a post-capitalist society.

This dissertation disputes immaterial labour theory. It does so through an analysis of work in the contemporary AI Industry – which by *post-operaismo*'s own definitions counts as immaterial labour. I argue that contemporary work in the AI Industry does not evince the qualities attributed to it by *post-operaismo*. On the contrary, work in the AI Industry suggests the continued relevance of Marx's original analyses of capitalism.

## Acknowledgments

This dissertation could never have come to completion without the supervisory influence of Dr. Nick Dyer-Witthford whose wide-ranging and continual acquisition of knowledge sets a high bar for an aspiring academic. It is hard to express my appreciation for the opportunities provided by my time working with him.

Second, I need to acknowledge the singular group of friends that I have been fortunate enough to accrue in the five years spent producing this dissertation. They are Tyler Wagner, Victoria O'Meara, Ryan Mack, Katie Oates, Chandell Gosse, Daniel Goyette, Joe Kim, Hannah Waitsches, Indranil Chakraborty, Ryan Schroeder and Yimin Chin.

I am particularly indebted to Dr. Vincent and Stacy Manzerolle who opened their home to me on my bedraggled arrival in London. I must also mention the particular importance of Dr. Atle Mikkola Kjøsén, who initiated me into the rites of value form analysis.

I also thank my committee members, Dr. Christian Fuchs, Dr. Anabel Quan-Haase, Dr. Alison Hearn and Dr. Chris Keep.

I must also thank Marc Lalonde, manager of the Grad Club, where I tended bar for most of the time spent writing this dissertation, for being a great boss.

Finally, I salute the Bruce Boys for years of enduring friendship. In addition to those already mentioned above, this includes Boh Pidskalny, Peter Riediger, Kyle Lefaive and Caleb Farrugia.

## Table of Contents

Abstract.....	ii
Summary for Lay Audience.....	iii
Acknowledgments.....	iv
List of Figures.....	ix
List of Appendices.....	x
Chapter 1.....	1
1 Introduction: The Artificial Intelligence Industry and Immaterial Labour.....	1
1.1 A Fourth Industrial Revolution?.....	1
1.2 Defining Artificial Intelligence.....	6
1.3 Actually-Existing AI.....	8
1.4 AI Industry Critical Literature Review.....	9
1.5 AI and Work: Selective Literature Review.....	13
1.6 Primary Research Question and Methodology.....	18
1.7 Chapter Outline.....	19
1.8 Contribution to Existing Body of Knowledge.....	22
Chapter 2.....	23
2 Labour, Capital, Machine.....	23
2.1 Classical Political Economy.....	24
2.2 Marx on Value and Labour.....	27
2.3 Marx on Machines: “devourers of living labour”.....	34
2.4 The Fragment on Machines.....	39
2.5 Marxism(s).....	41
2.6 Labour Process Theory.....	43
2.7 Theorists of Cybernetic Capitalism.....	47
2.8 From <i>operaismo</i> to <i>post-operaismo</i> .....	51

2.9	<i>Post-operaismo</i> .....	54
2.10	The New Reading of Marx (NRM).....	60
2.11	Conclusion .....	64
Chapter 3.....		65
3	A Political Economic History of the AI Industry.....	65
3.1	The Historical Context.....	65
3.2	The Advent of AI.....	71
3.3	The AI Winter .....	77
3.4	The Birth of the AI Industry: Expert Systems .....	79
3.5	Strategic Computing: AI and the State .....	83
3.6	The Decline of Expert Systems.....	87
3.7	The Rise of Machine Learning .....	90
3.8	Canada, Deep Learning and the contemporary AI Industry .....	94
3.9	Conclusion .....	100
Chapter 4.....		101
4	Political Economy of the AI Industry .....	101
4.1	Charting the AI Industry .....	101
4.2	AI Tech Giants .....	105
4.3	AI Dinosaurs .....	109
4.4	AI Startups .....	110
4.5	AI Think Tanks .....	111
4.6	AI and the Contemporary State.....	114
4.7	AI Industry Concentration .....	116
4.8	Open Source AI, Clouds, AI Chips.....	119
4.9	Labour in the AI Industry .....	121
4.10	Composition of AI Industry Labour .....	126

4.11 AI Industry Labour Organization.....	129
4.12 Conclusion .....	131
Chapter 5.....	132
5 The Production of Machine Learning: A Labour Process Analysis .....	132
5.1 Group Selection .....	133
5.2 Recruitment.....	133
5.3 Interview Preparation.....	134
5.4 Conducting the Interviews and Composition of Interviewees.....	135
5.5 Exploratory Interviews and Follow-up Research .....	137
5.6 The Machine Learning Labour Process .....	137
5.7 Data Processing.....	139
5.8 Model Building .....	140
5.9 Deployment.....	143
5.10 The Commodity Form of AI.....	144
5.11 Empirical Control of the Machine Learning Labour Process .....	149
5.12 AI is Automation.....	155
5.13 The Automation of AI Work.....	157
5.14 Automating the Machine Learning Labour Process .....	161
5.15 Automation Without Codification .....	163
5.16 Other Forms of Automation in Machine Learning .....	165
5.17 Conclusion .....	167
Chapter 6.....	168
6 The New Autonomy of Immaterial Labour .....	168
6.1 Extant Criticisms of Immaterial Labour Theory.....	169
6.2 The Technological Logic of New Autonomy .....	171
6.3 Human-Machine Hybridization .....	172



6.4 Abstract Cooperation .....	175
6.5 New Autonomy from Capital.....	178
6.6 AI Work and Immaterial Labour Theory.....	181
6.7 Autonomy of What?.....	185
6.8 Conclusion .....	191
Chapter 7.....	193
7 Conclusion .....	193
7.1 The Importance of Automation.....	193
7.2 Limitations and Ways to Address Them.....	194
7.3 Seize the Means of Cognition? .....	197
References.....	201
Appendices.....	247
Curriculum Vitae .....	256

## List of Figures

Figure 2: Waterfall Methodology (Royce 1970, 329) .....	150
Figure 3: Agile Methodology (Kuruppu 2019).....	151
Figure 4: Google AutoML (Lu 2019) .....	165

## List of Appendices

Appendix 1: Recruitment Blurb.....	247
Appendix 2: Recruitment Poster.....	248
Appendix 3: Letter of Information and Consent.....	249
Appendix 4 : Ethics Approval.....	252
Appendix 5 : Interview Questions .....	253

## Chapter 1

### 1 Introduction: The Artificial Intelligence Industry and Immaterial Labour

#### 1.1 A Fourth Industrial Revolution?

Karl Marx conducted his analysis of capitalism in the wake of what has come to be known as the “industrial revolution,” in which steam and water powered machine-based manufacturing became the predominant form of industry in Europe and parts of North America (Toynbee 1997).<sup>1</sup> The industrial revolution occurred somewhere between 1760 to 1840; Marx published the first volume of *Capital* in 1867. The first industrial revolution was followed by a second industrial revolution, characterized by the proliferation of electricity among other innovations, took place from around 1870 up to the First World War in 1914 (Landes 2003). This was followed by a third industrial revolution, characterized by the proliferation of digital networks and ICTs after the Second World War onto the present day (Rifkin 2011).

Today we are purportedly living in the midst of a fourth industrial revolution which began in the 2010s. According to Klaus Schwab (2017), founder and executive chairman of the World Economic Forum (WEF), the fourth industrial revolution is defined by a “confluence of emerging technology breakthroughs, covering wide-ranging fields such as artificial intelligence (AI), robotics, the internet of things (IoT), autonomous vehicles, 3D printing, nanotechnology, biotechnology, materials science, energy storage and quantum computing” among others (7). Similarly, the MIT economists Erik Brynjolfsson and

---

<sup>1</sup> While its use is widespread in both academic and popular culture, the notion of the industrial revolution has been justly criticized for oversimplifying complex social processes, attributing deterministic power to technology and overlooking the gradual nature of historical change (De Vries 1994; Žmolek 2013). Cameron (1982) laments the notion’s wide usage because it “has no scientific standing and conveys a grossly misleading impression of the nature of economic change” (377). Accepting these critiques fully, I use the concept here only in a schematic way, as shorthand means for summarizing technological milieus. I draw on the notions of successive industrial revolutions only to produce an overview of the technological evolution of capital, without committing to hard epochal boundaries or endorsing the arguments mounted by their promoters.

Andrew McAfee (2014) argue that after a long first machine age, which followed the first industrial revolution, we are only now entering a “second machine age” defined by the same array of technologies noted by Schwab. Brynjolfsson and McAfee (2017) suggest that of these technologies, AI will turn out to be the most significant because, like electricity or the combustion engine, it is a general-purpose technology with nearly limitless applications (3-4). Perhaps unsurprisingly, AI luminaries agree. AI expert and venture capitalist Andrew Ng has called AI the “new electricity” (Eckert 2016). He asserts that “[j]ust as electricity transformed almost everything 100 years ago, today I actually have a hard time thinking of an industry that I don’t think AI will transform in the next several years” (Lynch 2017). This sentiment is widespread. The consulting firm Accenture proclaims that AI is “the future of growth” (Purdy and Daugherty 2016). And while another such firm, Gartner (2018), puts deep learning AI at the apex of its hype cycle – “the peak of inflated expectations” – so far enthusiasm has yet to deflate.

Impressive research advances in AI continue to emerge. One AI system reportedly plays the classic computer game DOOM, remembers the environment spatially and temporally and has “dreams” about these memories. In its so-called dreams, the system reconstructs the game environment based on its experiences within it and autonomously improves its ability to play the game by exploring this dream world (Ha and Schmidhuber 2018).

Another AI system reportedly can diagnose skin cancer by visual inspection with accuracy comparable to human dermatologists (Esteva et al. 2017). Yet another (AlphaGo) defeats the world champion of Go, and is then itself defeated by its next iteration (AlphaGo Zero) – which received no human training at all (Silver et al. 2017).

Outside the research laboratory, self-driving trucks, powered by AI, now operate in limited areas, automated retail stores use AI to dispense with (some) retail employees and AI is widely used to conduct advanced analytics anywhere large quantities of data are available. In 2016, for the first time, the top five largest companies in the world by market value all came from the same country and industry – US big tech (Mosco 2017, 65). The same five were still on top in 2018: Apple (\$926.9B), Amazon (\$777.8B), Alphabet/Google (\$766.4B), Microsoft (\$750.6B), and Facebook (\$541.5B), with Chinese tech giant Alibaba in sixth place (\$499.4B) (Forbes 2018). Significantly, all of

these companies are now deeply involved in the research and production of AI. In 2017, Microsoft dropped its “Mobile First” mantra in favour of a new focus on AI (Darrow 2017). As I will show in the following chapters, a turbulent, burgeoning industry – the AI Industry – has formed around these high tech firms. I define the AI Industry as an industry composed of capitalist firms producing AI not only for their own purposes, but also for sale as commodities and means of production for other capitalists, as Chapters 3 and 4 elaborate.

Yet at the same time as the AI Industry expands, public figures including Elon Musk and Bill Gates as well as physicist Stephen Hawking have respectively described AI as “our biggest existential threat,” a “concern,” and a potential “end of the human race” (Gibbs 2014; Holly 2015). In slightly less eschatological terms, concerns have been widely voiced over the possibility of widespread technological unemployment driven by a new wave of automation powered by fourth industrial revolution technologies such as AI. While predictions vary in terms of when and how AI will impinge on labour, a vague consensus is emerging that, sooner or later, and in some as of yet unknown way, advanced capitalist societies are going to have to address the question of what to do when AI automates much of society’s labour.

For this reason, among others, the analysis conducted by Marx 200 years ago remains relevant. Marx argues that while surplus value derives from labour, machines are one of capital’s primary means for harvesting that value and controlling the humans on which it relies. I will suggest in Chapter 2 that Marx’s thought can be usefully schematized with the conceptual triad: labour, capital, machine. Marx shows how, driven by its need to harvest surplus value, capital has an inherent and ceaseless drive to technological revolution. Capital therefore tends towards an increasingly automated, machinic state.

Marx was not, however, flatly anti-technology. His work suggests rather that technology freed from the constraints of capitalist employment might be used for emancipatory ends and the general betterment of humanity. Subsequent Marxian thinkers have further developed, and sometimes revised, Marx’s framework to analyze the evolution of capitalism throughout the second and third revolutions. Marxist thinkers tend to call the

era wrought by the second industrial revolution “Fordism” which finds its most famous expression in the deskilled assembly line tending “mass worker” (Wright 2002). The digital third industrial revolution has been described, by Marxists and others, as a transition to “post-Fordism,” a disputed term which, at the minimum, refers to how capital sought to overcome to organized power of the mass worker through the deployment of information and communications technologies (ICTs) (Amin 1994, 1-34).

The strain of Marxian thought called *post-operaismo* (post workerism) has been particularly influential in the analysis of the Post-Fordist era. *Post-operaismo* is an offshoot of the *operaismo* (workerism) school of Marxism developed in Italy in the first half of the 20<sup>th</sup> century, which posited a “Copernican revolution” in Marxist thought by emphasizing the power of labour to direct the functioning of capital, rather than the other way around (Toscano 2009).<sup>2</sup> *Post-operaismo* thinkers argue that ICTs increase the autonomy and power of labour, transforming old forms of work into “immaterial labour” which capital cannot control (Lazzarato 1996; Hardt and Negri 2001). *Post-operaismo* claims that Marx’s approach is incapable of grasping this emergent situation because he lived in the age of steam-powered factories, not the internet. *Post-operaismo* posits a fundamental “change in the quality and nature of labor ... [wherein] information and communication have come to play a foundational role in production processes” (Hardt and Negri 2001, 289). All labour tends toward a “cyborg, high-tech form” (Dyer-Witheford 2005, 152-153).

The shift to immaterial labour is purported to have dramatic effects on the antagonistic relation between capital and labour. As Camfield (2007) puts it, immaterial labour is claimed to be of “world-historic importance” because it is “dissolving the division of time between work and non-work, creating a new commonality, undermining qualitative divisions among working people, producing life outside the sway of capital and making possible the popular unity of singularities that can achieve absolute democracy” (30).

---

<sup>2</sup> The English-speaking world has often referred to *operaismo* and its descendents, like *post-operaismo*, as “autonomist” Marxisms to emphasize their fundamental tenet that labour is autonomous from capital (Cleaver 1979; Witheford 1994).

Indeed, Antonio Negri, (1996) perhaps the most prominent figure in *post-operaismo*, once asked “[i]s this the third industrial revolution or the time of transition to communism?” (156).

While the internet never catapulted the world into communism, *post-operaismo* continues to advance its immaterial labour theory with little revision. Hardt and Negri (2017) evinces little theoretical difference from Hardt and Negri (2001). Many activists and academics continue to use immaterial labour theory as a theoretical framework (Kolođlugil 2015; Grizzioti 2018). Machines are still supposed to be increasing the power of immaterial labour contra capital, even as swaths of capital gather around a new nucleus in the AI Industry. One does not have to subscribe to Schwab’s whole fourth industrial revolution program to acknowledge that in the last 20 years the technological milieu has substantially changed. Machine learning AI, to take only one example, only emerged in its contemporary form around 2010. It is a fair question to ask, then, whether immaterial labour theory, conceived during the third industrial revolution, fares well amidst the fourth?<sup>3</sup>

This dissertation contests the validity and utility of *post-operaismo*’s immaterial labour theory. Workers in the AI Industry are immaterial labourers *par excellence*. Yet, if we examine this industry, we find that its workers do not possess the attributes accorded to immaterial labour. While *post-operaismo* makes numerous claims about immaterial labour, my focus is its claim for the increased autonomy of immaterial labour *vis a vis* capital.

My critique is founded on an empirical analysis of the AI Industry, which was produced through documentary analysis and interviews conducted with workers and management in the AI Industry. This analysis of the AI Industry was then interpreted via a reading of

---

<sup>3</sup> Left accelerationism (Williams and Srnicek 2014; Mason 2016; Bastani 2019), extolled by *post-operaismo* thinkers like Negri (2014), might arguably be considered a new *post-operaismo* for the fourth industrial revolution. Left accelerationism, however, inherits from *post-operaismo* its flaws, as my co-authors and I have argued (Dyer-Witheford, Kjøsen and Steinhoff 2019). An assessment of left accelerationism falls beyond the scope of this dissertation.



Marx influenced by the New Reading of Marx (NRM) school. NRM has a very different focus from *post-operaismo*, focusing on the continued relevance of Marx's theory of value rather than the disruptive significance of technological change. I suggest that work in the AI Industry does not exhibit the autonomy attributed to immaterial labour by *post-operaismo* and that NRM allows us to better understand labour in the AI Industry as subject to continued control by capital. Further, my analysis reveals that the AI Industry is giving birth to a new type of automation – automation without codification. I will argue that automation without codification signifies a possible technological change within capital whereby capital, not labour, appears to be increasing its autonomy. In short, my argument is that *post-operaismo* has it backwards. By the same token, however, the advance of AI also puts in sight a horizon over which the concept of value on which NRM rests itself threatens to vanish. This issue is taken up in the concluding chapter.

## 1.2 Defining Artificial Intelligence

Defining AI is difficult because there is no consensus on how to define intelligence nor on how it should be implemented in machines (Legg and Hutter 2007; Wang 2008). One early AI researcher defined AI as “making a machine behave in ways that would be called intelligent if a human were so behaving” (McCarthy et al. 1955). Another defines AI as “that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment” (Nilsson 2010, xiii). The problem of definition has been further complicated by a phenomenon, called the AI Effect, whereby once a problem in AI is solved, it is no longer regarded as requiring intelligence and is relegated to mere computation. McCorduck (2004) notes that “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’” (204).

This dissertation adopts a working definition of AI derived from Wang (2008) and Kaplan (2016). For Wang (2008), intelligence is “adaptation with insufficient knowledge and resources” which means that an intelligent system “is finite, works in real-time, is open to novel tasks, and learns from experience” (371). Kaplan (2016) defines the “essence of intelligence” as 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 behavior” (5-6). The definitions offered by Kaplan and Wang have the benefit of delimiting the field of what is possibly intelligent or artificially-intelligent by positing human intelligence as a tacit benchmark, with its criteria of temporal finitude, finite quantity of information and generality of application. These criteria allow us to differentiate between computation and AI.<sup>4</sup>

Another important definitional point concerns the distinction between AI and robotics.<sup>5</sup> The roboticist Winfield has offered 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).

Most importantly, robots have bodies. AI is software which runs on computing hardware. Advanced humanoid robots, such as Honda’s ASIMO, employ AI in perception and planning locomotion, but a robot body does not necessarily entail AI. Most of the existing robot population today is comprised of relatively dumb industrial robot arms, first deployed in 1951, though AI will likely increasingly suffuse robotics.<sup>6</sup>

---

<sup>4</sup> A program which can solve a difficult problem given 400,000,000 years of computation to conduct a brute force search through every possible solution does not possess intelligence in the sense we use when speaking of human intelligence, which operates with a much shorter timescale. Likewise, human intelligence functions with limited data all the time, except in constrained formal situations with “perfect information” (games such as checkers or chess) (Mycielski 1992, 42).

<sup>5</sup> AI should also not be associated inherently with consciousness. No existing AI systems have consciousness in any conventional sense of the term, excepting theories such as that of Koch (2012) in which consciousness is a property of matter in general. The question of machine consciousness is not relevant to this dissertation. The analysis presented here does not depend on machine consciousness being physically or even logically possible, nor does it depend on the impossibility of such.

<sup>6</sup> Research into humanoid care robots is ongoing in Japan and “cobots” or AI-powered collaborative robots are being developed to tackle tasks which machines alone cannot yet handle.

### 1.3 Actually-Existing AI

This dissertation is primarily concerned with actually-existing AI. This means that, other than some speculations in the conclusion, I am not concerned with fictional representations and possible future forms that AI might or might not take. The definitions of AI provided by Wang and Kaplan help distinguish actually-existing AI from speculative forms by plotting them on a continuum of intelligence. It is a handy heuristic: the more domains an AI can function in, and the less time, knowledge and resources it requires, the more intelligence it exhibits.

Actually-existing AI is called “narrow” AI because the “vast majority of current AI approaches ... are primarily designed to address narrow tasks” (Johnson, Hofmann, Hutton & Bignell 2016, 4246). All commercial applications of AI, and most AI research in general, have focused on such narrow task-based tools. Just as a word processor is not useful for generating 3D models, an AI system for recognizing faces is not going to predict stock prices. Narrow is thus a useful qualifier for actually-existing AI because it denotes that whatever intelligence is at work is of a qualitatively different sort than that of humans. Unless otherwise specified, all references to AI in this dissertation will refer to actually-existing narrow AI.

The technical literature on AI lies beyond the scope of this project, but the non-technical reader can find useful information in the introductory chapters of AI textbooks, such as the popular Russel and Norvig (2009). There also exist many non-technical introductory works covering the basics of AI (Warwick 2013; Frankish and Ramsey 2014; Boden 2016; Kaplan 2016) and machine learning in particular (Alpaydin 2016; Domingos 2015). There are several histories of AI (McCorduck 2004; Crevier 1993).<sup>7</sup> The most recent and most comprehensive history is Nilsson (2010). Despite its impressive clarity

---

<sup>7</sup> There are also histories of particular well-known AI systems and projects, such as Japan’s Fifth Generation computing programme and robotization efforts (Feigenbaum and McCorduck 1984; Shodt 1988), DARPA’s strategic computing initiative (Roland and Shiman 2002), IBM’s chess-playing Deep Blue (Hsu 2004) as well as IBM’s Jeopardy-playing Watson (Ferrucci 2012).

and scope, Nilsson’s history, like those that preceded it, focuses primarily on concepts and personalities and rarely on how AI became an industry in its own right.

Narrow AI may be contrasted with general intelligence, such as that of humans. Hypothetical future AI with human-like capacities is called artificial general intelligence (AGI) and “human-level AI” (Nilsson 2005). AGI is defined as AI with “the capacity for efficient cross-domain optimization” or “the ability to transfer learning from one domain to other domains” (Muehlhauser 2013). AGI is thus, theoretically, able to engage intelligently in a wide variety of contexts and to apply knowledge learned in one context to novel situations. By most assessments, the development of AGI, if it is possible, lies far in the future, although some futurists place its appearance as early as 2045 (Kurzweil 2005). While AGI research remains highly speculative, there are at least 45 active AGI projects across academia and industry (Baum 2017).<sup>8</sup>

## 1.4 AI Industry Critical Literature Review

AI has been subject to critical thought since Turing – arguably the first philosopher of AI – but there has been a resurgence of critical interest in AI since around 2010. While this dissertation draws primarily on Marxist literature, diverse other perspectives have considered the social, political and theoretical implications of AI. This section surveys several.

There is a diverse field of philosophical thought about AI. The philosophical literature on AI delves into classical philosophical fields such as philosophy of mind, metaphysics and epistemology and engages with sempiternal philosophical debates concerning agency, materialism and the nature of cognition and the self. Much of the philosophy of AI comes from the analytic philosophy tradition (Boden 1990; Copeland 1993). There has been no philosophical consensus on any aspect of AI. Some philosophers are dubious of AI’s

---

<sup>8</sup> AGI is exceeded in speculative amplitude by artificial superintelligence (ASI) or an AI “that greatly outperform[s] the best current human minds across many very general cognitive domains” (Bostrom 2014, 63). ASI is usually imagined as the outcome of a scenario in which AGI is invented and upgrades itself into an ASI with god-like powers. This now-classic science fiction scenario receives serious academic discussion at prestigious places including Oxford University’s Future of Humanity Institute.

philosophical assumptions and goals (Dreyfus 1979, 1992; Penrose 1989; Eklia 2008) while others are more measured (Haugeland 1989). Still others have argued that AI (particularly the artificial neural network) upsets the foundations of philosophical thought itself, calling for radical revisions (Churchland 1989, Churchland and Sejnowski 1994, Churchland 2013). A singular perspective has been developed by Negarestani (2015b, 2016, 2018) for whom the development of AI, particularly AGI, is the ultimate goal of philosophy.

There is also the field of machine ethics, which “is concerned with giving machines ethical principles, or a procedure for discovering a way to resolve the ethical dilemmas they might encounter” (Anderson and Leigh Anderson 2011, 1). Some machine ethics work focuses on the near term (White 2015; Wallach and Allen 2008) while other work is concerned with the possibilities and risks of possible future AI (Majot and Yampolskiy 2014; Müller 2016).

Another strand of philosophical thought derives from the continental philosophy tradition. While they have rarely spoken about AI specifically in much detail, continental thinkers have theorized machines more generally and considered their social and philosophical implications. Most of these thinkers draw on, to a greater or lesser degree, the work of Deleuze and Guattari (Johnston 2008; Pasquinelli 2015b). The famous “Manifesto for Cyborgs” (Haraway 1990) initiated a perspective which melds feminist continental philosophy with a critical, yet optimistic, appraisal of technology so as to break down categorical barriers between genders, species, humans and machines – developing what is now called a “posthuman” perspective (Hayles 2008; Braidotti 2013; Roden 2014). Technologies such as AI are often positioned in posthumanist theory as constitutive or potentially constitutive of new modes of being for humans and/or new ways of thinking about humans and machines. One suggestive example of this is Hayles (2017) who advances the notion of the “cognitive nonconscious,” by which cognition is to be distinguished from consciousness, and which thereby implies that nonconscious things can think (for a similar notion, see Harari 2016).

Posthumanism must be distinguished from another philosophical movement in which AI is also implicated: transhumanism (Ranisch and Sorgner 2014; Tegmark 2017).

Transhumanism is a cultural movement which advocates for the technological augmentation and/or replacement of humans with advanced technologies, including nanotechnology, robotics, biotechnologies and AI. Transhumanism was brought into mainstream discourse by inventor and futurist Kurzweil (1990, 2000, 2005) who elaborated on ideas drawn largely from the roboticist Moravec (1988, 2000).

Transhumanism was also popularized by More (1993), and was brought into academic discourse by Bostrom (2001) and others. Discussions about transhumanism often include mention of the Technological Singularity theory, proposed first by Good (1966) and named as such by Vinge (2013 [1993]). The Technological Singularity theory maintains some variation on the idea that once a sufficiently advanced AI is created (usually with approximately human level capabilities) it will quickly begin to improve itself. This AI, the story goes, will iteratively increase its powers, become an AGI, and eventually possess god-like powers (Shanahan 2015).<sup>9</sup>

More relevant to this dissertation is work on the social, political and ethical implications of actually-existing AI. This literature comes from diverse perspectives. Some of this literature concerns issues that arise when technologies with some degree of autonomy, such as AI, are integrated into weapons making “lethal autonomous systems” (Arkin 2009, 49) or “killer robots” (Scharre 2018; Krishnan 2016; Garcia 2014; Docherty 2012). There has also been interest in the ethical dimensions of civilian autonomous technologies, such as driverless cars (Goodall 2014; Bonnefon, Shariff and Rahwan 2016).

---

<sup>9</sup> The nature of such an ASI has been subject to debate. Kurzweil (2005) has suggested that humanity will be included in the ascendance of the AI while others such as Bostrom (2014) have called for cautionary preplanning due to the extreme danger presented by such a scenario. Overall, transhumanism has had a number of well-known detractors (Fukuyama 2004; Joy 2000; Barrat 2013). The lively debate over transhumanism has been well-documented in volumes of collected essays (Hansell 2011; More and Vita-More 2013).

There is also a rapidly expanding body of work which has been usefully dubbed “critical algorithm studies” (Social Media Collective 2016).<sup>10</sup> In their excellent review, Mittelstadt et al. (2016) divide this body of work into seven categories. Three are epistemic in nature. In my own words, they are: the difficulties in understanding output of algorithms (Holzinger 2018), the inscrutability of how algorithmic systems produce the output they do (Pasquale (2016), and the ruinous influence of bad training data (Mittelstadt et al. 2016, 4-5). The authors also discuss concerns around the attribution of responsibility to algorithms, which they call “traceability” (Mittelstadt et al. 2016, 5). While important to the critical study of AI, these are not of immediate interest to my dissertation. However, Mittelstadt et al. (2016) also point out two normative concerns that are of immediate interest: unfair outcomes and transformative effects.

Unfair outcomes refers to how the output of an algorithm may be assessed in terms of perceived fairness. Such work investigates how AI is implicated in social power relationships including gender, race and class and often exacerbates inequalities therein. The predominantly white male identity of AI creators is often the starting point of such analyses (Adam 2006; Crawford 2016). O’Neil (2017) demonstrates the wide variety of algorithmic discrimination in work, finance and other spheres of life. Noble (2018) and Larson, Angwin and Parris Jr. (2016) focus on race and gender, showing how racial biases of engineers can be baked into the AI systems they create. Eubanks (2018) focuses specifically on how such systems repeat historical patterns of oppressing the poor. Others have elaborated the complex intertwinings between AI, Silicon Valley, white supremacy and the alt-right. (Pein 2012; Golumbia 2019).

Transformative effects refers to the capacity for algorithms, and thus AI, to “affect how we conceptualize the world, and modify its social and political organization” (Mittelstadt et al. 2016, 5). A primary concern of this type of work is how AI-powered automation of work will affect societies. Some analyses argue that AI-powered automation is leading to

---

<sup>10</sup> Critical algorithm studies is diverse and includes work from disciplines ranging from communication studies (Granka 2010), human geography (Graham 2005), linguistic anthropology (Kockelman 2014), cultural studies (Striphas 2015; Mackenzie 2015) and media archaeology (Mackenzie 2017). Work has even been done exploring the theological dimensions of AI (Geraci 2012).

a narrowing of human cognitive abilities (Carr 2014; Danaher 2019) or a society-wide stupidity (Stiegler 2017).

In somewhat less theoretical vein, there is an ongoing debate amongst economists, computer scientists and business analysts about the effects of AI on employment. Some predict wide scale technological unemployment—a full scale “crisis of work” (Steiner 2012, Standage 2016, Ford 2009; 2015). Others believe job losses to AI-powered automation will be far smaller, and counterbalanced by employment gains due to overall economic growth and the emergence of new types of AI-related work (T. Lee 2018; Hawksworth, Berriman and Goel 2018). While a full assessment of this debate is beyond the scope of this dissertation, I review the basics in the next section. This dissertation instead focuses on the work of producing AI and whether *post-operaismo*’s theory of immaterial labour is able to adequately explain it.

## 1.5 AI and Work: Selective Literature Review

Researchers at Oxford found that 47% of all US jobs are potentially automatable through the application of “computerization” over the next ten to twenty years (Frey and Osborne 2013, 1). A later analysis from the OECD paints a much lighter picture with only 9% of jobs across OECD countries as automatable (Arntz, Gregory and Zierhan 2016, 4). The matter remains unsettled. Hawksworth, Berriman and Goel (2018) argue that by the 2030s around 30% of all jobs will be automatable, while analysts at McKinsey suggest that half of all paid activities in the global workforce have some potential for automation with currently existing technologies. While only less than 5% are completely automatable, they assert that for about 60% of all jobs 30% of tasks could be automated (Manyika et al. 2017).

The 2013 Oxford report focuses on jobs as a whole while the latter adopt finer-grained task-based perspectives, but even this, argues K.F. Lee (2018), is not sufficient to grasp the extent of AI’s impact on work. Both the job and task based approaches consider AI-powered automation as conducting a “one-to-one replacement of a machine for a human worker” but it is necessary to also consider how AI automation might also induce “ground-up disruptions” which decimate human employment in a particular industry by



completely reconfiguring its structure to deliver its core function in a totally different way (K.F. Lee 2018, 177-178).<sup>11</sup> Lee cites as an example the news app Toutiao which lacks the skills of an editor but still performs the curation of a newsfeed. This kind of AI automation, K.F. Lee (2018) suggests, could affect as much as 10% of the workforce in the USA (178).

No consensus can be derived from these reports. As Acemoglu and Restrepo (2018) note, “we are far from a satisfactory understanding of how automation in general, and AI and robotics in particular, impact the labor market and productivity” (1). Analysts from Brookings agree, holding that “the discourse appears to be ... suggesting that automation will bring neither apocalypse nor utopia, but instead both benefits and stress alike” (Muro, Maxim and Whiton 2019). This ambivalence is expressed throughout the literature by the repetition of the mantra that new technologies will destroy jobs through automation but at the same time enable the creation of new jobs (Susskind and Susskind 2015, 286; Agrawal, Gans and Goldfarb 2018, 223-4).

It is, regardless, important to note that the employment outcomes of AI-powered automation do not only depend on whether or not a job is “automatable”. The most basic Marxist, or even neoclassical, economics reminds us that for machines to be introduced they generally have to be not just technically feasible, in a laboratory setting, but also—at least over a foreseeable length of time—*cheaper* to the capitalist, in an industrial context, than the labour they replace. Since such large-scale machinic displacement of labour is likely to, all else being equal, increase the number of unemployed workers, and thus decrease the average cost of labour, it would simultaneously make humans more attractive to capital than machines. Therefore, it is likely that any enlargements in technological unemployment from AI introduction will unfold in a very staggered and circuitous process, rather than an abrupt vaporization of jobs.

---

<sup>11</sup> Here Lee neatly describes the Marxian notion of real subsumption – the reconfiguration of production processes to forms adequate to the needs of capital (Marx 1990, 501). I discuss this in the next chapter.

Another mantra that is repeated throughout the AI and work literature is what I call centaur theory.<sup>12</sup> In 1997, the chess champion Garry Kasparov lost to IBM's AI Deep Blue. He was the first human champion to lose to an AI. He later went on to develop a style of chess in which human-machine teams (centaurs) would work together. The best centaurs could defeat both human masters and the best AI systems (Case 2018). Centaur theory is a popular solution to the uncertain future of work in the face of AI-based automation.

Brynjolfsson and McAfee (2014) is one of the most influential works on the economic implications of AI as well as centaur theory. Brynjolfsson and McAfee (2014) assert that the emergence of “real, useful artificial intelligence” is one of the “most important one-time events in our history” (90). AI and widespread digital networks are “more important than anything since the Industrial Revolution” (Brynjolfsson and McAfee 2014, 90). In their assessment, as AI proliferates “costs will go down, outcomes will improve, and our lives will get better,” even if low-skill, routine jobs are largely eliminated (Brynjolfsson and McAfee 2014, 91). Davenport and Kirby (2016) similarly argue that in the era of automation of cognitive work, “the parts of our jobs we’ll keep are just the parts that can’t be codified” (14). The notion is that we need develop AI such that it augments, rather than automates, human capabilities (see also Markoff 2016; O’Reilly 2017).

Exactly how labour is divided in the human-AI centaur varies from author to author, but the general point is the same: humans will do whatever cannot be automated yet. For Brynjolfsson and McAfee (2014), this is creative thinking: “people who are good at idea creation will continue to have a comparative advantage” against machines (192).

Agrawal, Gans and Goldfarb (2018) recommend that machines perform routine data-heavy predictions and humans use judgment to make predictions in irregular, low-data scenarios (68-9). They hold that as machine learning-powered prediction becomes cheaper, the value of and demand for, non-automatable human judgment will increase

---

<sup>12</sup> A genealogy of centaur theory, which there is no space to conduct here, would likely begin with the work of Douglas Engelbart who advanced the notion of using technology for human “intelligence augmentation” (see Bardini 2000).

(Agrawal, Gans and Goldfarb 2018, 178). Daugherty and Wilson (2018) call for a similar model in which AI does repetitive tasks while humans employ their judgment. The creation of centaurs will, they suggest, create so many jobs that technological unemployment should not be worried about (5-8). Frank, Roehrig and Pring (2017) call the centaur a “knowledge economy exoskeleton” (151) and assert that “for the vast majority of professions, [AI] will actually enhance and protect employment” (8). Microsoft (2018) promotes a similarly rosy theory of “human-centered AI” (136). Schenker (2017) holds that professional jobs will be largely immune to automation for the foreseeable future (63).

However, some centaur theories are less rosy. Susskind and Susskind (2015) argue that even the non-routine and creative jobs of professionals will be done by “increasingly capable machines, operating on their own, or with non-specialist users” (159). The one sphere where humans might persist, they suggest, is in positions involving moral responsibility (Susskind and Susskind 2015, 279-284). It is unclear how many jobs this particular specialization will provide. McAfee and Brynjolfsson (2017) have refined their original centaur theory, three years later, in light of advances in ML. While they originally reserved creative thinking as safely reserved for human workers, the authors now note that “machines are getting quite good at coming up with powerful new ideas on their own” (McAfee and Brynjolfsson 2017, 111). They are:

confident that the ability to work effectively with people’s emotional states and social drives will remain a deeply human skill for some time to come. This implies a novel way to combine minds and machines as we move deeper into the second machine age: let the computers take the lead on making decisions (or judgments, predictions, diagnoses, and so on), then let people take the lead if others need to be convinced or persuaded to go along with these decisions (McAfee and Brynjolfsson 2017, 123).

Their new conclusion is that human workers now need to specialize in affective labour. In sum, whether blindingly optimistic or more guarded, centaur theory holds that AI can or should be deployed as an augmentation, rather than automation, technology. However,

centaur theorists tend not to discuss how the economic system of capitalism compels the development and deployment of technologies such that any line between what machines can and cannot do is unlikely to remain fixed for long.

Some analyses of AI and work have paid more attention to essential dynamics of capital. Kaplan (2015) expects that AI-powered automation will exacerbate persistent social problems like inequality and unemployment (3). He even asserts that widespread AI automation will reveal that Marx was right that capitalism “is a losing proposition for workers” (Kaplan 2015, 11). He suggests that AI presents a dire picture of the future that Marx could not imagine: a society in which capital no longer needs labour: “[t]he real problem is that the wealthy will need few, if any, people to work for them at all” (Kaplan 2015, 11). His solution to the problems that will be caused by AI automation is a framework of “free-market solutions to address the underlying structural problems we are creating” (Kaplan 2015, 13). These include a new financial instrument, the “job mortgage” as well as a “government-certified measure of corporate ownership ... [called] the public benefit index” (Kaplan 2015, 14). The problem with Kaplan’s work is that while he recognizes that AI presents a heightening of the labour-capital conflict, he does not recognize this conflict as irreducible under capitalism.

Ford (2015) develops a clearer picture of the situation because he grasps the essential dynamics of the process of capital. Like Kaplan, he holds that a new wave of AI-enabled automation makes the elimination of routine and predictable jobs certain and that this will have disastrous consequences for labour. Ubiquitous computing power and the “distributed machine intelligence that accompanies it” means that any new fields of work created will not be labour-intensive (Ford 2015, 176). Unlike Kaplan, he does not think the problems posed by AI can be addressed without modifying the capitalist mode of production itself. Ford is skeptical of the centaur theory argument and also of free market fixes like those of Kaplan, because as he correctly notes, automation is not the individual choice of management or individuals: “progression toward ever more automation is not an artifact of ‘design philosophy’ or the personal preferences of engineers: it is fundamentally driven by capitalism” (Ford 2015, 255-256). He therefore argues that stopping the tide of AI automation “would require modifying the basic incentives built

into the market economy” (Ford 2015, 256). Ford stops short of advocating for revolution, instead opting for the establishment of universal basic income (UBI). This presents its own difficulties, which are beyond the scope of this dissertation.<sup>13</sup>

## 1.6 Primary Research Question and Methodology

This dissertation provides an answer to the following primary question:

*Does work in the AI Industry evince the new autonomy from capital attributed to immaterial labour by post-operaismo?*

By answering this question, this dissertation aims to contribute to a better understanding of contemporary high-tech labour. Adapting from Raniero Panzieri (1965), it aims to “defy all kinds of mystical ideas” about labour by testing *post-operaismo*’s theoretical claims with the concrete example of work in the AI Industry. To do so, it deploys a multi-pronged methodology. This dissertation combines documentary analysis, theoretical analysis and qualitative interviews with AI workers and management. Documentary analysis and interviews were used to produce a history, political economy analysis and labour process analysis of the AI Industry. Theoretical analysis was then applied to assess *post-operaismo*’s claim for a new autonomy of immaterial labour. It is important to explain why this analysis required three different (history, political economy and labour process) analyses of the AI Industry.

This dissertation studies the AI labour process because this is where the purported new autonomy of immaterial labour should be evident. Marxist analyses have argued that technological change in the labour process is one of capital’s main ways of increasing the exploitation of labour and thus its harvest of surplus value. As Wright (1981) puts it:

The Marxist model ... directs our research efforts toward those transformations of the socio-technical conditions of production that directly impinge on surplus labour. It is for this reason that the Marxist

---

<sup>13</sup> For a Marxian analysis of UBI, see Dyer-Witthford, Kjøsén and Steinhoff (2019).

analysis of production revolves around the analysis of the labour process (63).

Yet, analysis of the labour process alone is not sufficient. The labour process must be situated within a larger political economic context to be properly understood. Labour processes are structured and changed according to the dynamics of the antagonism between labour and capital and the industry at large. One must attempt to elaborate the “relationship between what goes on in the workplace and the forms it assumes and is determined in and by virtue of the market” (Pitts 2018b, loc 2008). In other words, a labour process study should not neglect what Marxists call the valorization process, or the process by which capital augments itself as value.<sup>14</sup> Further, since any given industry, with its particular political economic dynamics, does not arise spontaneously, but is a product of the social and material situation that came before it, an attempt at historicizing is also necessary. Indeed, for a Marxist, which is necessarily a materialist, analysis, the “development of the economic formation of society is viewed as a process of natural history” (Marx 1990, 92). This dissertation thus situates today’s AI Industry against a historical backdrop to show where it came from and how it came to be.

## 1.7 Chapter Outline

The second chapter offers a survey of several different schools of Marxist theory. This survey is conducted through the lens of a conceptual triad formed by three concepts: labour, capital and machine. It shows how these three concepts, although always connected in Marxist thought, have been articulated in different ways. The chapter also shows that despite their interest in technology, Marxist theorists have rarely discussed AI or the work of producing AI. The chapter begins by briefly discussing the conceptual milieu in which Marx’s work emerged – classical political economy. Then it moves on to the elaborate Marx’s critique of political economy and his theorization of capital’s inherent tendency towards an increasingly machinic or automated state. It shows how, for

---

<sup>14</sup> Wright (1981) notes that “[w]hile the labour theory of value has been implicit in most of this [labour process] research, rarely is it explicitly into the conceptualization of the problem” (65).

Marx, machines are capital's favored weapon in its ongoing struggle against the human labour it relies on. It also elaborates how Marx's theory of capital is essentially recursive. Recursion is a motif that will reappear throughout this dissertation.

Next, a variety of more recent Marxisms are discussed, with attention to those theorists who have attempted to apply Marxism to cybernetic capitalism. Labour process theory, which focuses on concrete labour practices and how capital attempts to wrest control of them from labour via the introduction of machines, is also elaborated. The significant inversion of the Marxist theory of machines conducted by *post-operaismo*, with its theory of immaterial labour, is then discussed. *Post-operaismo*'s claim for a new autonomy of immaterial labour from capital – the object of critique for this dissertation – is introduced. Finally, the New Reading of Marx (NRM) is introduced. This approach, which is critical of *post-operaismo*, is positioned as an alternative way to understand AI and AI work.

The third chapter consists of a political economic history of the AI Industry. AI is situated in a broader history of automation and computer technologies for which recursion – simply put, the application of computing to computer development – is a central property. The history of the AI Industry begins in the 1950s with the advent of AI research in the USA and Britain and develops through the “AI winter” into its first manifestation around GOFAI in the 1980s. Then, after advances in machine learning, the contemporary AI Industry first appears around 2010 and rapidly expands in the following years.

The fourth chapter conducts a political economy analysis of the contemporary AI Industry, drawing on both interview data and documentary research. It considers the capital side of the AI industry, including its scale and scope, the various types of companies that make it up, the types of products produced and the peculiar dynamics that distinguish it from others, including the proliferation of open source AI tools. It also considers the labour side of the industry. It analyzes the hierarchy of AI labour, the massive salaries earned in the industry, the widespread sexism and racism of the industry, the shadow world of online microworkers engaged in cleaning and moderating data, as

well as the recent emergence of organization and activism in AI and other high tech labour, which has been historically characterized by its apoliticism.

In the fifth chapter, I switch from macro scale analysis to focus on the labour process of the contemporary AI Industry. In particular, this chapter examines how machine learning AI is produced. It draws extensively on interview data as well as documentary research, so I present a full account of the interview methodology at the start of this chapter. The chapter describes the three technical stages of the machine learning labour process and then discusses four key themes drawn from interview data with people working in the AI Industry. These are the commodity form of AI, empirical control of the machine learning labour process, AI as an automation technology and the automation of AI work. The automation of AI work with the technology of automated machine learning (AutoML) is then detailed for all three stages of the machine learning labour process. I argue that AutoML represents a new type of automation, which I call automation without codification.

The sixth chapter draws on the previous chapters to mount a critique of *post-operaismo*'s purported new autonomy of immaterial labour. It first attempts to reconstruct, in optimal form, *post-operaismo*'s technological argument for new autonomy. It then argues, that based on the analysis of the AI Industry presented in this dissertation, AI work – which is immaterial labour *par excellence* – does not exhibit in any sense the new autonomy described by *post-operaismo*. It argues that, on the contrary, AI work, with the emerging technology of AutoML, instead seems to indicate a trajectory for the increasing autonomy of capital, rather than labour. I argue that NRM, with its insistence on the continued relevance of Marx's theory of value, is better equipped to grasp this scenario. However, even NRM is not fully prepared to deal with the possibilities presented by the long-term development of AI-powered automation, which, I suggest, could undermine NRM's concept of the interdependence of human labour and value. Thus, neither *post-operaismo* nor NRM emerge unscathed from the encounter with AI staged in this dissertation.



The seventh chapter argues that while the analysis of AI work and the critique of *post-operaismo* mounted here may seem cripplingly negative, it indicates one way forward for critical thought and practice, beyond the chronic pessimism of NRM, even if the path suggested is far from the breezy autonomy posited by *post-operaismo*. This chapter also discusses limitations of the study.

## 1.8 Contribution to Existing Body of Knowledge

This dissertation offers two original contributions to the existing body of knowledge. One of these contributions derives from its Marxist engagement with AI. As Chapter 2 shows, few Marxist thinkers of any stripe have engaged with contemporary AI in any detailed or systemic manner. One exception to this is Dyer-Witthford, Kjøsén and Steinhoff (2019), of which I am co-author. This dissertation differs in focus from that monograph, however, in that it draws on interviews with people working in the AI Industry and that it explores the emerging technology of AutoML, which has significant ramifications for Marxist theory. The second contribution offered by this dissertation concerns the empirically-grounded critique mounted against *post-operaismo*'s claims for a new autonomy of immaterial labour. While *post-operaismo* has received critiques on numerous fronts, as Chapter 2 and 6 show, its claims for a new autonomy have received minor attention, and have not been addressed on the specifically technological level on which they are formulated.

In sum, this dissertation contributes to the sophistication of critical theory by offering a Marxist approach to studying AI and by demonstrating the invalidity of a fundamental premise of a very popular theoretical framework. In so doing, it aims to contribute to a better understanding of cybernetic capitalism.

## Chapter 2

### 2 Labour, Capital, Machine

Marxist thinkers have devoted little attention to the analysis of AI, but they have, in various ways, described how technology and labour are intertwined in the process of capital valorization. This chapter does not survey every one of these numerous approaches, but focuses on those that have been most energetically developed (or which I intend to develop) in the analysis of digital technology and digital labour, and its most recent manifestation, AI.<sup>15</sup> In order to develop a better Marxist approach to AI, this chapter surveys how several schools of Marxist thought have theorized the entanglement of labour and technology under capitalism, forming a conceptual triad: labour, capital, machine. I follow this triad from its first manifestations in classical political economy, through Marx's critique of political economy, and on to its diverse permutations in several more recent Marxian schools of thought which have attempted to grapple with increasingly computerized or cybernetic capitalism. These include labour process theory, the New Reading of Marx, *operaismo*, and *post-operaismo* (the object of critique of this dissertation).<sup>16</sup> While I assess *post-operaismo*'s particular claim for a new autonomy of immaterial labour in detail in Chapter 6, this chapter first maps the Marxist theoretical landscape more generally, so that *post-operaismo* may be distinguished from rival schools of thought.

---

<sup>15</sup> For example, this dissertation does not pursue the highly abstruse debates about the tendency of the rate of profit to fall due to an increase in the organic composition of capital. This is because, important as such analysis is, it has not to date yielded any specific in-depth examination of AI or of the labour involved in creating it

<sup>16</sup> A note is necessary on an unavoidable issue of circularity. The reader should be aware that since I endorse a NRM interpretation of Marx (and there is no interpretation-free version of Marx), my exposition of his work necessarily slants towards a NRM reading. While I have tried to quote minimal commentators in my exposition of Marx, some were too useful to pass up.

## 2.1 Classical Political Economy

The diverse school of Marxist thought have their roots in the work of the 18<sup>th</sup> and 19<sup>th</sup> century political economists who pioneered the notion that “the wealth of nations is founded on the productive power of labour” (Bonfeld 2014, 22). While the young Marx focused his critique philosophers such as Hegel and Feuerbach, the mature Marx shifted his attention to the political economy of Adam Smith and David Ricardo, among others, who sought to understand the capitalist system, which was transitioning to industrial production, around them. Marx’s mature work *Capital* is subtitled “a critique of political economy”. Yet, while Marx did severely critique the political economists, his work built upon the foundation they established. Along with the centrality of labour, the political economists also recognized the fundamental importance of machines. And they discerned, in different ways, the development of an antagonism between labour and machines under capitalism – an antagonism inherently connected to the production of value.

Adam Smith was perhaps the first to recognize the integral link between technology and labour in capitalism. Labour was central to Smith’s theory because it was the labour expended in production that determined the value of a commodity and therefore the basis of capitalist economies. Smith held that: “The value of any commodity ... to the person who possesses it, and who means not to use or consume it himself, but to exchange it for other commodities, is equal to the quantity of labour which it enables him to purchase or command. Labour ... is the real measure of the exchangeable value of all commodities” (Smith 1991 [1776], 36).

By positing labour as the source of value, Smith initiated a process of, according to Pitts (2018b), “plunging deeper into the layers below” surface phenomena such as exchange (29-30). Pitts (2018b) holds that Smith significantly looked beyond particular physical acts of labour to labour’s “generalized role in the production process as a whole” and to the social form taken by labour in a society founded on exchange (28). He supposed that value derived not only from labour, but also from the “work” of the other two classes: landowner (rent) and capitalist (interest) as well, thus expressing a vague notion of value as comprised from a social relation – something that Marx would later elaborate. But

although Smith initiated this critical gesture, his concept of value remains expressed in price and thus lacks the critical force of Marx's concept of value, as I will show below.

Smith is also notable, contrary to how he is referenced by mainstream economists today, for his acknowledgement that capital and labour possess inherently opposed interests. Toussaint (2009) even argues that Smith adumbrated the theory of class struggle described in Marx and Engels' *Communist Manifesto* seventy years after the publication of Smith's opus *The Wealth of Nations*. Smith notes that:

The workmen desire to get as much, the masters to give as little as possible. The former are disposed to combine in order to raise, the latter in order to lower the wages of labour. It is not, however, difficult to foresee which of the two parties must, upon all ordinary occasions, have the advantage in the dispute, and force the other into a compliance with their terms (Smith 1991 [1776], 98).

Smith also noted that capital was driven towards the use of technology in production, even in his pre-industrial time. He famously detailed and extolled the division of labour in manufacturing for its enhancement of productivity of labour through specialization and his respect for machines runs along similar lines. Machines are valuable because they "facilitate and abridge labour, and enable one man to do the work of many" (Smith 1991 [1776], 21-22). Therefore, a "much smaller quantity of labour becomes requisite for executing any particular piece of work" (Smith 1991 [1776], 338). Smith, however, did not connect the antagonism between classes to machinery.

Such positive evaluations of technology were widely held in the years following the publication of Smith's *The Wealth of Nations*. Two proponents of this type of discourse frequently cited by Marx are Charles Babbage and Andrew Ure. Babbage (1832) extolled machinery for its ability to "supersede the skill and power of the human arm". He described three advantages derived from the use of machines: "The addition which they make to human power. The economy they produce of human time. The conversion of substances apparently common and worthless into valuable products" (Babbage 1832). Ure (1835) went so far as to assert that: "The constant aim and effect of scientific

improvement in manufactures are philanthropic, as they tend to relieve the workmen either from niceties of adjustment which exhaust his mind and fatigue his eyes, or from painful repetition of effort which distort or wear out his frame” (8). Interestingly, both Ure and Babbage were not only promoters of industrial machinery, but also, in different ways, early forayers into computing (Zimmerman 1997). Babbage designed an early mechanical calculator or proto-computer called the Difference Engine while Ure’s musing on manufacturing might be considered proto-cybernetic in how they interlink human and machine.

David Ricardo complicated the rosy view on machines. Like Smith, Ricardo advocated a labour theory of value. He held that the: “value of a commodity, or the quantity of any other commodity for which it will exchange, depends on the relative quantity of labour which is necessary for its production, and not on the greater or less compensation which is paid for that labour” (Ricardo 2001 [1821], 8). Ricardo thus shifted focus from of price to embodied labour as the locus of value. He dismissed Smith’s ruminations on the social nature of value in favour of a theory of value founded on the concrete act of labour at the point of production (Pitts 2018b, 29-30). This embodied conception of labour allowed him to conceptualize machines “and other aids to production” as “accumulated labour” (Dinerstein and Neary 2002, 14-15, cited in Pitts 2018b, 30). Here Ricardo connects the products of labour back to labour – a point later developed by Marx.

At first, Ricardo saw machines as Smith, Babbage and Ure did: benefitting both capital and labour. However, by the third edition of his *On the Principles of Political Economy and Taxation*, Ricardo changed his optimistic tune (Kurz 2010, 1197-1198). He then argued that while the introduction of machines into production certainly benefits the capitalist because it increases output and decreases labour costs, “the substitution of machinery for human labour, is often very injurious to the interests of the class of labourers” (Ricardo 2001 [1821], 283). Ricardo thus affirms what Smith noted, that different classes in society may have differing and contrary interests, but he also connects class antagonism to the application of technology in the workplace. The “same cause which may increase the net revenue of the country, may at the same time render the population redundant, and deteriorate the condition of the labourer” (Ricardo 2001

[1821], 284). Ricardo even speculated on the possibility of a fully-automated mode of production, in what we might call Ricardo's Fragment on Machines: "If machinery could do all the work that labour now does, there would be no demand for labour. Nobody would be entitled to consume any thing who was not a capitalist, and who could not buy or hire a machine" (Ricardo 1951–1973, VIII: 399–400, cited in Kurz 2010, 1195). Here Ricardo envisions a capitalism in which machines are functionally-identical to humans and wholly replace them; humanity as a whole becomes what Marx would later call a "surplus population" (Marx 1990, 517). Marx would also later develop the ramifications of machines taking on increasingly more capacities once reserved for humans, as the next section shows.

In sum, classical political economy set the scene for Marxist theory by triangulating labour, capital and machines. As a whole, however, political economy only managed to describe the functioning of capitalism from its surface. In the words of Pitts (2018b), while Ricardo attempted to dive below surface appearances of capital to unearth the truth of labour, "he did not pose the question as to why and how products of labour become value-bearing commodities on this basis [of labour]. It is this question that forms the springboard of Marx's analysis" (30). In other words, while political economy asks about how capitalism works, Marx asks why it is that capitalism works in the way it does.

## 2.2 Marx on Value and Labour

Marx's critique of political economy draws on both Smith and Ricardo, but mainly, and crucially, in order to reveal the inadequacies of their theorizations of capital. Marx's position is elaborated throughout his mature work: the three volumes of *Capital* (Marx 1990, 1992, 1991) and the notes he made while conceiving *Capital*, now called *Grundrisse* (Marx 1993). Marx deepens the contradiction that Ricardo noted between labour and machines under capital, and posits a fundamental antagonism between the two. But he also paints a picture, though never completed, that suggests that machines might ultimately betray capital to the benefit of labour. Different interpretations of the properties of machinery in Marx are one reason for the theoretical divergence amongst different schools of Marxism.

To grasp the significance of Marx's critique of political economy one has to see how his use of the categories of political economy (such as labour) differs from that of the political economists. As indicated above, Marx's goal is to get below the surface appearances of capital. However, this does not mean that reading Marx's gospel will dispel the false preconceptions of the reader and somehow give her an undistorted access to reality. This is because the surface appearances of capital are nonetheless real things with real effects. As Smith (2009) puts it: "Marx's main theoretical task is to explain how the social relations of capitalism necessarily generate appearances that distort what capital essentially is (appearances that none the less have material effects)" (125). One ready way to grasp this is by looking at Marx's concept of labour.

Like Smith and Ricardo, Marx holds that value is comprised of labour. Labour is, first of all, a generic term for workers and the working class. However, labour is also an activity universally performed by humans. Labour "is a condition of human existence which is independent of all forms of society; it is an eternal human necessity which mediates the metabolism between man and nature, and therefore life itself" (Marx 1990, 132). Under capitalism, however, labour takes the particular form of wage labour, in which a capitalist pays a worker a wage for control over the worker's capacity to labour for a given period of time. For Marx, unlike Smith and Ricardo, wage labour exists in two forms: concrete and abstract. Concrete labour is the human capacity to process the environment into novel useful forms, creating use-values, or useful products or services. This is why Marx deems labour a transhistorical universal. However, at the same time, concrete labour always refers to a particular act of labour: writing a program, drilling a fencepost or assembling a circuit board.

However, within the capitalist mode of production there exists a second form of labour, which Marx calls "abstract labour". At the same time as any given act of concrete labour consists of a particular kind of labour, it is also "an expenditure of human labour-power" (Marx 1990, 137). Labour-power is the capacity to labour or "the aggregate of those mental and physical capabilities existing in the physical form, the living personality, of a human being" (Marx 1990, 270). Marx (1990) is clear on this: "When we speak of capacity for labour, we do not speak of labour" (277). One might assume then that

abstract labour refers simply to all labour in a broad, generic sense. But this is not the case. Marx (1990) states that abstract labour has the function of forming the value of commodities, in particular its “quality of being equal” does so (137). But how could various acts of abstract labour which supervene on particular acts of concrete labour, as the application of a set of capacities for labouring, be measured and judged equal? By duration (Marx 1990, 164). Since duration is a pure quantity, value is thus defined as “congealed quantities of homogenous [abstract] human labour” measured temporally (Marx 1990, 128). It is only because of this homogeneity that one commodity can be compared with another in terms of value, and thus exchanged. Value as congealed labour is what makes commodity exchange possible: “[t]he equality of the kinds of human labour takes on a physical form in the equal objectivity of the products of labour as values” (Marx 1990, 164).

Yet the actual quantity of hours worked cannot determine value alone, or merely working slowly would increase the value of products. Therefore, Marx (1990) reasons, the value of a commodity is determined by the average “[s]ocially necessary labour-time” for its production (129). This refers to the “labour-time required to produce any use-value under the conditions of production normal for a given society and with the average degree of skill and intensity of labour prevalent in that society” (Marx 1990, 129). The value of a commodity is thus not equivalent to the amount of time spent working on it by its producer, but rather the average time that it takes to produce such a commodity in that particular time and place. Socially necessary labour time thus reveals value, which appears as a thing or property of things, as a complex social relation. Marx (1990) states that value’s “objective character ... is ... purely social” (138-9). It is not, contra Ricardo, imparted to the commodity by the concrete labour of producing it, nor is it some substance that inheres in commodities after their production: “[n]ot an atom of matter enters into the objectivity of commodities as values” (Marx 1990, 138). Value arises from the social form of commodity production.

Value is central to Marx’s analysis, but like labour, it manifests in two ways. Marx distinguishes exchange value (which is expressed as price) from value, which is determined by socially necessary labour time. The two do not necessarily coincide. In



fact, Marx argues that capital is not primarily driven to maximize exchange value. Instead, the aim of capital is to produce value, but “not just value ... also surplus-value” (Marx 1990, 293). Surplus-value is a quantity of value which capital appropriates from labour without recompense. This appropriation is the overarching goal which immanently configures the capitalist mode of production. How does capital appropriate surplus-value from labour? Against today’s common sense, Marx holds that surplus-value does not arise from selling commodities for more than they were produced for (this generates profit, which is not the same). Rather, surplus-value arises from the unequal exchange between capital and labour.

To grasp Marx’s argument one must realize that the labour-power of the worker is a commodity which the worker sells to the capitalist (Marx 1990, 271). The value of labour-power is determined by the average socially-necessary cost for its production, like any other commodity. This takes the form of things the worker needs for the “reproduction of himself and his maintenance” such as food and shelter (Marx 1990, 274). The time it takes a worker to earn enough money to successfully reproduce her labour-power is defined by Marx as “necessary labour” (1990, 325). The capitalist buys a temporal quantity of labour-power and aims, in general, to pay as little as possible without impinging on the reproduction of labour-power. This is because any time the worker works beyond necessary labour is “surplus labour” in which the worker produces value for which the capitalist has not paid an equivalent – surplus value (Marx 1990, 325). The commodity of labour-power is unique in that it produces more value than is paid for it. This unequal exchange defines Marx’s technical term exploitation and it explains why, for Marx (1990), there is an inherent antagonism between labour and capital (418-421).

What then is capital? Marx (1990) defines the “general form of capital” as  $M-C-M'$  “where  $M' = \dots$  the original sum advanced plus an increment ... This increment or excess over the original value I call surplus-value” (251). This formula represents the process by which money ( $M$ ) is invested in the production of commodities ( $C$ ) which are sold and turned into more money than was initially invested ( $M'$ ). This formula demonstrates that while surplus-value is appropriated during production, it is not realized

until capital exchanges the commodities produced for money. Thus to say that surplus-value is extracted by capital by the exploitation of labour in production is only part true. The circulation of commodities is also essential. If labour is successfully exploited, and if capital successfully realizes the value of its commodities, it now has more value than when it hired the worker – this capital has accumulated value and has thus been “valorized” (Marx 1990, 252). This valorized capital can then be invested in new circuits of commodity production. Thus Marx (1990) asserts that “capital is not a thing” but rather, a process (953). Capital “only becomes real capital, value valorizing itself, in the course of the process” (Marx 1990, 1061). And not only is capital not a thing, it is not finite: “the circulation of money as capital is an end in itself ... The movement of capital is therefore limitless” (Marx 1990, 253). This conception of capital as an infinite, self-referential process is perhaps Marx’s greatest innovation because it reveals capital as, most simply, the blind compulsion towards the accumulation of surplus-value. As Nick Land (2017) puts it, capital “appeals to nothing beyond itself, it is inherently nihilistic. It has no conceivable meaning beside self-amplification. It grows in order to grow”. Capital is an infinite, self-referential process, and Marx suggests, it can achieve an increasingly automatic operation. Marx describes capital as “the subject of a process in which, constantly assuming the form in turn of money and commodities, it changes its own magnitude, throws off surplus value from itself ... and valorizes itself independently” (1990, 255). Marx thus calls capital the “self-valorization” (1990, 255) of value, the “automatic subject” (1990, 255) and the “dominant subject” (1990, 255).

This is another way of understanding Marx’s assertion that “the characters who appear on the economic stage are merely personifications of economic relations; it is as the bearers of these economic relations that they come into contact with each other” (Marx 1990, 179). Humans, capitalist as well as worker, are actors directed by the script of capital’s prime directive, the increase of value: “The buyer of labour-power is nothing but the personification of objectified labour which cedes a part of itself to the worker in the form of the means of subsistence in order to annex the living labour-power for the benefit of the remaining portion, so as to keep itself intact and even to grow beyond its original size by virtue of this annexation” (Marx 1990, 1003-1004). To flesh out this position it is useful to turn to the concept of fetishism, which Marx supplies in his analysis of the

commodity. Marx suggests that to grasp how the commodity functions in capitalism we must make a comparison to the “misty realm of religion” where:

the products of the human brain appear as autonomous figures endowed with a life of their own, which enter into relations both with each other and the human race. So it is in the world of commodities and the products of men’s hands. I call this the fetishism which attaches itself to the products of labour as soon as they are produced as commodities, and is therefore inseparable from the production of commodities (Marx 1990, 165).

Marx argues “as soon as men start to work for each other in any way, their labour ... assumes a social form” (1990, 164). The social form of labour influences the form the product of labour assumes. The capitalist mode of production entails that products of labour take the form of commodities, that is, products intended for exchange. Because the wants and needs of life are commoditized, nearly all aspects of social life are constituted as exchange relations. The form the commodity takes:

reflects the social characteristics of men’s own labour as objective characteristics of the products of labour themselves, as the socio-natural properties of these things. Hence it also reflects the social relation of the producers to the sum total of labour as a social relation between objects, a relation which exists apart from and outside the producers (Marx 1990, 164-165).

Thus, under capital, “the definite social relation between men themselves ... assumes here, for them, the fantastic form of a relation between things” (Marx 1990, 165). The commodity obscures its own origin in labour and appears as value.

It is important to note how exactly Marx phrases the relation of humans to fetishism. Marx writes “the products of the human brain appear as autonomous figures” (1990, 165) and assume “for them, [humans] the fantastic form of...” (1990, 165). This indicates that there is an illusory character to this appearance. However, it is not that simple. Capitalists

are actually constrained to pursue valorization or face bankruptcy while workers are obliged to sell their labour-power or starve to death. The fetish-character of the commodity is both real and efficacious in the world as well as an illusion; it is a “real illusion” (Holloway 2002, 71), “real abstraction” (Toscano 2008) or “spectral objectivity” (Heinrich 2012, 52). It has objective effects on the world, but its ontological status is that of a relation among humans. In other terms, the social relations humans come to engage in eventually come to control them by appearing as things. Appearance thus means, for Marx, “a concrete social reality created on the basis of a mystified and disguised process” (Dyer-Witheford, Kjøsen and Steinhoff 2019, 21). Marx (1990) makes this clear:

Men do not therefore bring the products of their labour into relation with each other as values because they see these objects merely as the material integuments of homogeneous human labour. The reverse is true: by equating their different products to each other in exchange as values, they equate their different kinds of labour as human labour. They do this without being aware of it. Value ... does not have its description branded on its forehead; it rather transforms every product of labour into a social hieroglyphic. Later on, men try to decipher the hieroglyphic, to get behind the secret of their own social product: for the characteristic which objects of utility have of being values is as much men's social product as is their language (166-167).

The commodity and capital thus really do have “a life of their own” (Marx 1990, 165) at the same time that that life is comprised of the social relations of humans. Capital is a subject or “higher-order alien power” (Smith 2009, 123) even while it is a process comprised of humans socially relating in a particular way. This duality means that fetishism is not merely something in the mind of people. It is a material state of affairs and thus “cannot be overcome in thought alone” (Holloway 2002, 71).

In sum, Marx’s critique of political economy generates an image of capital as a potentially infinite process of the valorization of value, which through its functioning, systemically occludes its true nature through the generation of efficacious real illusions,

such as value. Valorization requires the exploitation of labour to capture surplus value, but capital appears as productive of value itself – and it is, insofar as value takes an objectified form which is recursively applied to labour, from whence it came. However, production, and thus valorization, depend on more than labour-power, they require also the means of production, including raw materials, facilities, tools, and machines.

To see how machines fit into Marx's view of capital, we now need to switch perspectives from the "valorization process" to the "labour process" (Marx 1990, 283). Since value is a non-physical real abstraction, the valorization process necessarily supervenes on the physical substratum of the concrete production of use-values. As Marx (1990) puts it: "Just as the commodity itself is a unity formed of use-value and value, so the process of production must be a unity, composed of the labour process and the process of creating value" (293). While the valorization process requires a labour process on which to supervene, the labour process is simultaneously determined in its nature by the exigencies of valorization. The machine is capital's preferred device by which to determine the nature of the labour process.

### 2.3 Marx on Machines: "devourers of living labour"<sup>17</sup>

While I follow Marxian tradition and speak of capital as such, it is important to recollect that capital is comprised of many individual capitals competing with one another on the market, as well as individually, against their more or less recalcitrant labour forces. The immanent drive of capital to valorization is thus inextricable from the competitive dynamic amongst individual capitals: "the immanent laws of capitalist production ... assert themselves as the coercive laws of competition" (Marx 1990, 433). Individual capitals thus employ a variety of techniques for minimizing necessary labour time and thus increasing surplus-value extraction.

The simplest way to do this is to lengthen the working day; the necessary labour stays the same while the surplus increases. This Marx calls "absolute surplus-value" (1990, 432).

---

<sup>17</sup> (Marx 1990, 983).

He also describes “relative surplus-value” (1990, 432) which is produced when a relative decrease in necessary labour time is achieved, with a consequent increase in surplus labour time. While the absolute approach “maxes out” due to the finite length of the day, the increase in relative surplus-value has no internal limit and thus is the primary axis of intercapitalist competition. Since the amount of necessary labour-time is determined by the cost of the reproduction of labour-power, the first can be decreased by reducing the cost of the second. Due to the need to capture relative surplus-value, “[c]apital therefore has an immanent drive, and a constant tendency, towards increasing the productivity of labour, in order to cheapen commodities and, by cheapening commodities, to cheapen the worker himself” (Marx 1990, 436-437).

Marx discusses several methods by which individual capitals compete to optimize the capture of relative surplus-value, including cooperation and the division of labour (Marx 1990, 452-456). While these remain essential to capital to this day, the most effective technique in capital’s repertoire, in Marx’s time as it is now, is machinery. Marx opens his chapter on machines by stating that the alleviation of the burden of work is “by no means the aim of the application of machinery under capitalism” (1990, 492). On the contrary, he argues that “[l]ike every other instrument for increasing the productivity of labour, machinery is intended to cheapen commodities and, by shortening the part of the working day in which the worker works for himself, to lengthen the ... part he gives to the capitalist for nothing (Marx 1990, 492). The machine “is a means for producing surplus-value” (Marx 1990, 492) not because it produces value itself, but because it allows the reduction of necessary labour-time: “[t]he productivity of the machine is therefore measured by the human labour-power it replaces” (Marx 1990, 513).

In contrast to variable capital or the wage, machines employed in capitalist production take the form of fixed capital. Fixed capital is bought and does not have labour-power which it might sell for a wage used to buy reproductive necessities. For Marx, labour-power is an exclusively human attribute. Machines cannot possess it. It exists “only as the capacity of the living individual” (Marx 1990, 274). Fixed capital cannot be exploited and cannot generate surplus value: “[m]achinery ... creates no new value but yields up its

own value to the product it serves to beget” (Marx 1990, 509). The machine “never adds more value than it loses, on an average, by depreciation” (Marx 1990, 509).

What exactly is a machine? Marx distinguishes machines from tools. He defines a machine as comprised of three parts: a motor mechanism, a transmission mechanism and a “tool or working mechanism” (1990, 494). When “the tool proper is taken from man and fitted into a mechanism, a machine takes the place of a mere implement” (1990, 495). This is because the human body is limited in terms of how many tools it can wield, the type of tools it can wield, and the forces it can exert. Machines have no necessary form and therefore no such “organic limitations” (Marx 1990, 495). Machines may also be linked together to form a “complex system of machinery” (Marx 1990, 499) in which an array of interconnected machines undergo a new division of labour (Marx 1990, 501). Such a system “constitutes ... a vast automaton as soon as it is driven by a self-acting prime mover” (Marx 1990, 502) even if parts of the machine need to be handled like a tool by a worker (Marx 1990, 502-503).

Such systems are impressive to Marx, but he notes a categorical difference only in what he calls “automatic” machines: “As soon as a machine executes, without man’s help, all the movements required to elaborate the raw material, and needs only supplementary assistance from the worker, we have an automatic system of machinery, capable of constant improvement in its details” (1990, 503). In an “organized system of machines to which motion is communicated by the transmitting mechanism from an automatic centre” Marx saw “the most developed form of production by machinery” (1990, 503). From the simple machine evolves a “mechanical monster whose body fills whole factories, and whose demonic power, at first hidden by the slow and measured motions of its gigantic members, finally bursts forth in the fast and feverish whirl of its countless working organs” (Marx 1990, 503). Here Marx describes the process of automation, *avant la lettre*.

Since, as a general rule, the “tendency and the result of the capitalist mode of production is steadily to increase the productivity of labour” (Marx 1990, 959), driven by its hunger for relative surplus-value, capital follows a trajectory towards an increasingly automatic

and machinic state. This Marx also called the increasingly “organic composition of capital,” defined by the relative increase of fixed as opposed to variable capital (Marx 1990, 762). This is why Marx refers to machinery as “capital’s material mode of existence” and the “material foundation of the capitalist mode of production” (Marx 1990, 554). Marx (1993) asserts: “[t]he development of the means of labour into machinery is not an accidental moment of capital, but is rather the historical reshaping of the traditional, inherited means of labour into a form adequate to capital” (694). This process of increasingly automatic machine production forms the labour process equivalent of automatic self-valorization in the valorization process.

Machinery’s adequacy to capital is best explained by reference to Marx’s notion of subsumption. There are two kinds of subsumption. Formal subsumption is simply when an existing labour process comes under the control of a capitalist and thus becomes part of a circuit of capital valorization (Marx 1990, 1019). This form of capitalist control does not alter the labour process itself, and, as such, can only produce absolute surplus value by lengthening the working day (Marx 1990, 1021). Competition, however, compels capitals to increase the production of relative surplus value by reconfiguring the entire labour process, primarily through the introduction of machines, such that “a specifically capitalist form of production comes into being (at the technological level too)” (Marx 1990, 1024, emphasis original). The assembly line is an easy example of this. Later Marxian thinkers have extended the concept of subsumption beyond production, as we will see in the following sections. The “essential difference,” says Marx, is that with real subsumption, via the introduction of machines, “[t]he worker has been appropriated by the process; but the process had previously to be adapted to the worker” (1990, 501). Marx noted three consequences for workers that had already occurred with the application of machines in production: allowing women and children in the workforce (therefore increasing the total number of workers and driving wages down) (1990, 517), lengthening the working day (1990, 526), and intensifying work itself (1990, 533). Yet more significant than these are the consequences entailed by the structure of capitalist production itself.



Marx holds that capital's drive towards the machinic entails an attack on labour as such. While the "division of labour develops ... labour-power in a one-sided way, by reducing it to the highly particularized skill of handling a special tool," a worse scenario obtains "[w]hen it becomes the job of the machine to handle this tool, [because] the use-value of the worker's labour-power vanishes, and with it its exchange-value" (Marx 1990, 557). In other words, "[t]he instrument of labour, when it takes the form of a machine, immediately becomes a competitor of the worker himself" (1990, 557). Thus, the "instrument of labour strikes down the worker" (1990, 559). The adversarial nature of the exploitation relationship is intensified into "complete and total antagonism with the advent of machinery" (Marx 1990, 558). Machinery is "the most powerful weapon for suppressing strikes, those periodic revolts of the working class against the autocracy of capital" (1990, 562). Marx even suggests that it "would be possible to write a whole history of the inventions made since 1830 for the sole purpose of providing capital with weapons against working-class revolt" (1990, 563).

By combining the perspectives of both labour process and valorization process, we can see that the logic of fetishism also applies to machines when they take the form of capital. As I have shown, according to Marx, value and thus capital, are comprised by dead labour. The "social character of his labour confronts the worker as something not merely alien, but hostile and antagonistic; when it appears before him objectified and personified in capital" (Marx 1990, 1025). This takes a particularly vivid form in machines, in which value as dead labour manifests concretely within the labour process, in material entities which exhibit the capacity to move and do work (though not labour) – today, we can add the capacity to think. In confronting machines, the worker is confronting his own creations, which have been captured by capital and turned against their creator: "[t]he entire development of the productive forces of socialized labour ... and together with it the use of science (the general product of social development), in the immediate process of production, takes the form of the productive power of capital" (Marx 1990, 1024, emphasis original). It is through machines, as well as money, that capital assumes the fetish character of autonomous productive power. Marx puts it poetically: "(P)ast labour – in the automaton and the machinery moved by it – steps forth as acting apparently in independence of [living] labour, it subordinates labour instead of being subordinate to it,

it is the iron man confronting the man of flesh and blood” (Marx 1861-63c, 30 quoted in Smith 2009, 125). The Marxist understanding of machinery thus seems clear. It is capital’s best weapon in its ongoing struggle to control and exploit labour for the production of surplus-value. This understanding is complicated, however, by a passage from Marx’s notes for Capital.

## 2.4 The Fragment on Machines

Marx expresses a different relation between labour, capital and machines in the Grundrisse. Here he imagines a highly automated future capitalism and draws from this scenario some startling conclusions. Marx asserts that the “full development of capital” occurs:

only when the means of labour has not only taken the economic form of fixed capital, but has also been suspended in its immediate form, and when fixed capital appears as a machine within the production process, opposite labour; and the entire production process appears as not subsumed under the direct skillfulness of the worker, but rather as the technological application of science (Marx 1993, 699).

Capital only achieves maturity when it becomes predominantly machinic. Direct labour becomes an “indispensable, but subordinate moment, compared to general scientific labour, technological application of natural sciences ... and to the general productive force arising from social combination” (Marx 1993, 700). In this scenario the “human being comes to relate more as watchman and regulator to the production process” (Marx 1993, 705). A few workers set in motion, and then oversee, and perhaps perform maintenance on, a vast system of machinery. The collective knowledge of the society (i.e. science), as represented in machines, is the most significant component of production. In this speculative future context, Marx deploys the peculiar term “general intellect”:

The development of fixed capital indicates to what degree general social knowledge has become a direct force of production, and to what degree, hence, the conditions of the process of social life itself have come under

the control of the general intellect and been transformed in accordance with it (Marx 1993, 706).

Social knowledge becomes a direct force of production, but only (contrary to the reading of *post-operaismo*, as I show in Chapter 6) once it has been implemented in machinery—because only in an object external to the human body can human skills and knowledge be called social—and most machinery in production functions as fixed capital. The general intellect refers to this ever-increasing manifold of skill and knowledge possessed by capital as machinery. Prior to its machinic implementation, the skill and knowledge of the social individual is called by Marx the “social brain” (1993, 694). Capital excises and machinically emulates aspects of the social brain and adds them to the general intellect.

As production grows increasingly automated, “the creation of real wealth comes to depend less on labour time ... than on the power of the agencies set in motion during labour time, whose ‘powerful effectiveness’ is ... out of all proportion to the direct labour time spent on their production” (Marx 1993, 704-705). However, since machines cannot be exploited and therefore cannot produce surplus value, as the general intellect expands and capital approaches full automation, there is less and less labour to exploit.

Eventually, capital cuts off its source of sustenance and starves itself to death. For this reason, Marx describes capital as “the moving contradiction” because “it presses to reduce labour time to a minimum, while it posits labour time, on the other side, as sole measure and source of wealth. Hence it diminishes labour time in the necessary form so as to increase it in the superfluous form” (1993, 706). While capital can intensify the exploitation of its remaining workers, Marx speculates that eventually the process will terminate in the collapse of capitalism. In effect, capital “works towards its own dissolution as the form dominating production” (Marx 1993, 700) and “production based on exchange breaks down” (Marx 1993, 705). Machinery betrays capital and destroys the conditions for its existence. The machinery developed by capital for maximum surplus value extraction can then be repurposed towards communist ends such as “the general reduction of the necessary labour of society to a minimum, which ... corresponds to the artistic, scientific etc. development of the individuals in the time set free” (Marx 1993, 06).

Although “The Fragment on Machines” is often described as a complete outlier in Marx’s writings, its vision of a capital that automates itself out of existence can be understood as an extreme version of a fundamental thematic that winds throughout his works. This is the idea that at a certain point of development within any mode of production, the “forces of production” – of which technology is a salient component – come into contradiction with the “relations of production” – property relations – causing revolutionary crises that result in the emergence of a new socio-economic order (Marx 1973). Nonetheless, the spectacular account of capital’s demise at the hands of its own technological creations is notable for its divergence with the more sober tone of *Capital*. And since it is a rather cryptic and fragmentary text which was never published in Marx’s lifetime, it is not possible to determine with certainty whether it really was a prediction of the future, some sort of speculative thought experiment, a parodic rendering of a position under critique or something else entirely. Given this uncertainty, reliance on “The Fragment” is highly controversial within contemporary Marxism – a point I return to later.

## 2.5 Marxism(s)

A number of guides to reading *Capital* have been written, advocating substantially different understandings of Marx (Harvey 2010; Heinrich 2012; Cleaver 2000 [1979]). Indeed, due to the diverse ways in which Marx has been read and applied, any use of Marx’s theory must make clear its adherence to a particular school (or schools) of Marxism. My research draws upon two Marxist schools of thought: the New Reading of Marx (NRM) and labour process theory. In this section, I will distinguish these from both orthodox Marxism and the influential paradigms of *operaismo* and *post-operaismo*. I will show how differently the triad labour, capital, machine has been configured by these various interpreters of Marx.

Marx only lived to see the publication of Volume 1 of *Capital*. The first dominant school of Marxism was driven by Engels’ editing, publication and interpretation of Marx’s work after his death (see Engels 1947), as well as the influential work of Karl Kautsky (1903). For Engels and Kautsky, Marx’s achievement was to have achieved a “scientific” understanding of the historical laws which govern social development (Elbe 2013). What is now often called orthodox Marxism led, via Lenin and the Bolshevik Revolution of

1917, to the official Communist Party doctrine of Marxism-Leninism promulgated by Joseph Stalin and Abram Deborin. Heinrich (2012) holds that “[i]deas in general circulation today concerning Marx and Marxian theory – whether these are appraised positively or negatively – are essentially based upon” this orthodox reading of Marx, which he calls “worldview Marxism” (26). Engels’ influence on this interpretation of Marx’s work was so substantial that Elbe (2013) suggests that orthodox Marxism should rather be called “Engelsism”. The orthodox reading of Marx posits a teleological character to social development, with communism standing at the end of history as the necessary resolution of the contradictions of capitalism. Capitalism must and will fail and communism must take its place; thus “humanity is subordinated to a ‘scientifically verifiable’ automatism of liberation” and Marx’s critical analysis is thus reduced to a “vulgar empiricism and historicism” (Elbe 2013).

The outcome of this was Communist Party doctrine in which technology was regarded as a capitalist development which could (and would) be appropriated by communist revolutionaries in the transition to a new mode of social being. While Lenin (1972 [1914]) first described Taylorism as “Man’s Enslavement by the Machine” he would, a few years later, strikingly assert that “the Taylor system, properly controlled and intelligently applied by the working people themselves, will serve as a reliable means of further reducing the obligatory [in the context of labor conscription] working day for the entire working population” (Lenin 1968 Vol. 42, 80, cited in Scoville 2001, 621). This technological optimism is also visible in Lenin’s (1920) famous dictum “Communism is soviet power + electrification of the whole country”.

A different understanding of Marx was initiated after the First World War by György Lukács (1972 [1923]) and Karl Korsch (2013 [1923]). According to Elbe (2013), Lukács was the first Marxist theorist to question whether Engels’ presentation of Marx reflected the intentions of its author. Lukács disputed the transhistorical laws and teleological formulations of orthodox Marxism and instead focused on the notion of ideology and the historically specific nature of capitalism. Other prominent theorists of what Elbe (2013) calls Western Marxism were Antonio Gramsci and the Frankfurt School critical theorists Theodor Adorno, Herbert Marcuse and Max Horkheimer. In different ways, these

thinkers were concerned with subjectivity and social praxis and often incorporated elements of psychoanalysis and Hegelian philosophy in addition to Marx's work. Frankfurt School thinkers rejected Lenin's appraisal of the prospects for unifying technology and labour. Marcuse's (2013 [1964]) *One Dimensional Man* is emblematic of this position, critiquing both capitalist and state socialist societies for their one-sided deployment of technological rationality and concomitant reduction of human potentialities to the service of narrow economic goals. Gramsci (1992 [1971]) held a more ambivalent view, recognizing a necessity, for communism, of the Taylorist rationalization of production, while objecting to its militarization in the service of War Communism (309-310).

In the wake of Western Marxism there emerged many influential schools of Marxian thought. The following sections examine several that made technology central to their analyses.

## 2.6 Labour Process Theory

Harry Braverman (1998) popularized labour process theory with his study of the theory of scientific management (or Taylorism) as propounded by Frederick Winslow Taylor in the late 19<sup>th</sup> and early 20<sup>th</sup> century. Labour process analysis draws on Marxist theory but eschews looking at the valorization process or the large-scale dynamics of capitalist labour and technology. Instead, labour process theory focuses on the particular concrete actions and procedures that comprise the labour processes of workers and the ways in which capital seeks to reconfigure these for its own ends through management practices, including the application of automation technologies. Braverman (1998) defines the labour process as "in general a process for creating useful values" which is also "specifically a process for the expansion of capital, the creation of a profit" (36). That is, the labour process, as Marx noted, is always subject to the exigencies of the valorization process.

Labour process theory thus draws inspiration from Marx's detailed analyses of the introduction of machines into 19<sup>th</sup> century factories in Chapter 15 of *Capital* Volume I. Braverman shows how Taylorist management subjected the movements and knowledge

of industrial workers to fine-grained scrutiny and documentation to make both available for analysis, and thus control, by capital via management. Taylorist analysis allowed capital to achieve new levels of real subsumption, increasing the division of labour and simplification of work. This had the consequence of deskilling labour, reducing the value of labour-power, and eventually, enabling the augmentation and replacement of labour by machines.

According to Braverman (1998), Taylorist deskilling is guided by three principles. The first is “the dissociation of the labor process from the skills of the workers” (78). This is the process by which management studies or interrogates workers to obtain their knowledge about their work tasks, so as to capture it and to control how it gets passed on to future workers. The second principle is “the separation of conception from execution” (Braverman 1998, 79). According to this principle, all creative and planning aspects of work should become the exclusive property of management. The third principle is “the use of [the] monopoly over knowledge to control each step of the labor process and its mode of execution” (Braverman 1998, 82). This means that the worker’s work should be prescribed, and followed mechanically, at as fine a grain as possible. Braverman (1998) argues that these principles could and would be applied to all kinds of work and that they remained “fundamental to all advanced work design or industrial engineering today” (77). Indeed, he holds that:

[m]odern management came into being on the basis of these principles ... Its role was to render conscious and systematic, the formerly unconscious tendency of capitalist production. It was to ensure that ... the worker would sink to the level of general and undifferentiated labor power, adaptable to a large range of simple tasks, while as science grew, it would be concentrated in the hands of management (Braverman 1998, 83).

With Taylorism, capital recognizes its own inherent drive towards extracting relative surplus-value and subjects it to systemic formalization, study and improvement. Labour becomes subject to all kinds of means of control, but most notably, machines. Machinery

“offers to management the opportunity to do by wholly mechanical means that which it had previously attempted to do by organizational and disciplinary means” (Braverman 1998, 134). Machines do this through “the progressive elimination of the control functions of the worker, insofar as possible, and their transfer to a device which is controlled, again insofar as possible, by management from outside the direct process” (Braverman 1998, 146). Witnessing the advent of information technology, in the form of numerical control, Braverman was pessimistic about the future because control would no longer be limited by the particular form of the machine. Instead, he expected a new “universality of the machine” which could be put to many uses without loss of control “since that control is no longer dependent upon its specialized internal construction” (Braverman 1998, 132). Braverman (1998) showed that, in the 1970s, clerical and so-called white collar office work was already subject to Taylorist management and suggested that the future would see the extension of such practices to “draftsmen and technicians, engineers and accountants, nurses and teachers” (282). He also noted how early programming work was almost from its inception subjected to a division of labour and suggested that it too could be Taylorized (Braverman 1998, 227-8).

Labour process theory received substantial criticism for a purported neglect to considerations of worker subjectivity, favouring instead structural relations of control. This conflict, and others, have been documented in several collections (Knights and Wilmott 1990; Wardell, Steiger and Meiksins 1999; Thompson and Smith 2010). Notable for labour process analyses which clearly include consideration of worker subjectivity is the work of David Noble (1979; 2011) for whom the labour process is recognized as a result of persistent worker resistance as well as capitalist control. Labour process theory was also criticized for its central deskilling thesis, even by those working within the tradition (Wood 1982). However, the issue is far from settled. As the editors of a recent volume on labour process theory and digital work note, new technologies and their applications for automation have brought the question of the degradation of work back into popular discourse (Briken, Chillias and Krywdinski 2017, 2).

Since Braverman, the labour process approach has been applied to various diverse fields of work including: machine shops (Noble 1984), telework, creative and knowledge work



(Thompson and Warhurst 1998; Huws 2003, 2014; McKinlay and Smith 2009), the work of Mobile app developers (Bergvall-Kåreborn and Howcroft 2013), airplane cabin crews (Taylor and Moore 2015), subcontracted service work (Grimshaw, Cartright, Johnson and Rubery 2017). Labour process analysis has also been contextualized within global supply chain networks (Newsome, Taylor, Bair and Rainnie 2015) and China, broadly (Liu and Smith 2016). However, very little labour process analysis has been done on the production of AI. The broad strokes of such an analysis are given in Dyer-Witford, Kjøsen and Steinhoff (2019) and will be elaborated in the fifth chapter of this dissertation.

In the near total absence of labour process analysis of AI production, my research has been informed by labour process analyses of software work in general. Kraft (1977, 1979) was the first to conduct a labour process analysis of software workers. Kraft (1977) argued that programmers were not exempt from the deskilling process outlined by Braverman (53). Kraft and Dubnoff (1986) concluded that “software work replicates rather than revolutionizes traditional relationships between managers and managed” (184). Cooley (1981) held that “Taylorism is destined not just for those who use computers or are in middle management, but for a range of intellectual ‘workers’ including scientists and technologists (51). Much early labour process work on software development echoed these sentiments (Cooley 1980; Duncan 1981; Orlikowski 1988; Greenbaum 1979, 1998; Ensmenger and Aspray 2000). Other voices dissented, arguing that there was no evidence of deskilling in software work (Tarallo 1987; Ainspain 1999; Hounshell 2000).

Rowena Barret has presented a more complicated picture, arguing that while traditional forms of direct control may not be evident in software development, these labourers are subject to control via “‘responsible autonomy’ which attempts to ‘harness the adaptability of labour power by giving workers leeway and encouraging them to adapt to changing situations in a manner beneficial to the firm’” (Friedman 1977, 78 cited in Barrett 2005, 79). Rasmussen and Johansen (2005) agree that “autonomy can be used as a strategy for increased control over the workers ... by devolving responsibility” (118). One particular manifestation of responsible autonomy is what Barret (2005) calls “technical autonomy”

or “management allowing software developers autonomy to develop the ‘best’ program using their skills and expertise” (82).

Yet, even as they identify these more subtle forms of control, labour process theorists have argued that it is “difficult to equate the practices and processes in software startups with an industrial assembly line. At every stage human rather than machine intervention predominates ... each project requires fresh planning and decisions. This reality stands in sharp contrast to the ‘one-best-way’ of Taylorized work settings” (Andrews, Lair and Landry 2005, 66; see also Barrett 2005). This is because software production is “dependent upon the skills of individuals and the synchronization of their disparate efforts” in ways management cannot implement from above (Andrews, Lair and Landry 2005, 59). Thus, a popular position from labour process theory today appears to be that software production is “craft rather than ... technical-oriented” and not to be at threat of deskilling because to “fully standardize computer programming ... would require the seemingly omniscient knowledge of both the emergent problems and the associated solutions” (Andrews, Lair and Landry 2005, 67). According to this view, software work thus admits of a fundamental shift in the triad labour, capital, machine: in the software industry capital cannot capture labour’s skills and knowledge and implement them in machines. The production of information machines is resistant to machinic implementation. Such a view, I will show, is shared by *post-operaismo*.<sup>18</sup>

## 2.7 Theorists of Cybernetic Capitalism

Before examining *post-operaismo*, I will survey some of the numerous Marxist and Marx-influenced thinkers who have attempted in diverse ways to apply and/or adapt

---

<sup>18</sup> Not all LP theorists agree with the conclusions of *post-operaismo*. Some explicitly reject them. Thompson and Briken (2017) argue that cognitive capitalism theory (an alternate name for *post-operaismo*) “is all too typical of the sweeping generalisations, unrepresentative exemplary industries and absence of plausible evidence that characterizes much of the recent social theorizing about capitalism” (Thompson and Briken 2017, 246). Further, they charge that “[t]heorists of cognitive capitalism ... significantly challenge LPA as they deny the idea of material production as a privileged site for the value extraction and antagonism between capital and labour” (Thompson and Briken 2017, 242).

Marx's triangulation of labour, capital and machine to the changed situation of capitalism since Marx wrote *Capital* between 1867-1883. Since technology is such a fundamental aspect of Marxist thought and almost all Marxist works engage with it to some degree, a complete survey is impossible here. Instead, I focus on the period of "cybernetic capitalism" (Freyre 1966, 106; Robins and Webster 1988; Peters 2015, 40). While there is no consensus definition for this term, I use it here to refer to the period of capitalism characterized by the proliferation of ICTs (and AI) after World War II (see Dyer-Witford, Kjosen and Steinhoff 2019, 51).

Marxian political economists of communication and media such as Schiller (1999; 2014) and Manzerolle and Kjosen (2012) have studied the processes by which capital became digitally networked and the substantial military origins of networking technologies such as the internet. Huws (2003; 2014) and others (Mosco and Wasko 1988; Mosco and McKercher 2008; Burston, Dyer-Witford and Hearn 2010) have explored the novel forms that labour has taken with the advent of digitally networked capital, including, early on, telework and later forms of internet work. Fuchs and Dyer-Witford (2013) argue for the utility of the Marxist approach to internet studies as well as digital labour across the spectrum of the digital industries (see also Fuchs 2014, Dyer-Witford 1999; 2015, Fuchs and Mosco 2015). Terranova (2004) initiated research on the "free labour" performed by users of internet sites such as social media, while others discerned the rise of a second new type of work often called the "gig economy" or on-demand economy, which involves work which is paid but sourced through digital platforms on an irregular basis (De Stefano 2015; Graham et al., 2017).

Some scholars have conducted critical analyses of the platform as both technology and business model (Gillespie 2010, Dyer-Witford 2013, Srnicek 2017). Many software and AI-producing companies operate a platform model. Mosco (2015) has studied the technology of cloud computing which on which platforms operate, as well as the digital high-tech industries more generally, although he accords minimal coverage to AI (Mosco 2017). Daubs and Manzerolle (2017) collects studies of mobile and ubiquitous media while Miller and Matviyenko (2014) collects some of the earliest studies of apps.

Three thinkers are especially relevant for configuring Marx's triangulation of labour, capital and machine in the age of AI. In different ways, the three of them grapple with the emerging capacities for recursion wrought by capital suffused with computers and other information technologies.

One is Morris-Suzuki (1984), who was among the first Marxists to think seriously about robots and AI. She holds that the "separation of hardware from software ... may be seen as constituting a revolutionary fission of the labour process itself" (1984, 112). The emergence of software, she holds, redefines completely the nature of machinery in two ways: "firstly that a single machine may be made to vary its movement without alteration to its mechanical structure ... secondly ... the worker's knowledge may be separated from the physical body of the worker and may itself become a commodity" (1984, 113). The worker's knowledge can then be represented as digital information, which Morris-Suzuki suggested capitalism would come to reorganize itself around: "information ... which contributes to productive processes—will become a commodity churned out by corporate enterprises almost as routinely and monotonously as cars flowing from an assembly lines" (1984, 114-115). The industrialization of information production entails that "it too becomes subject to the forces of automation" (1984, 120). With this, capital achieves a recursive capacity which empowers its technologies, leading Morris-Suzuki to argue that Braverman's deskilling thesis would indeed extend to intellectual labour with only the most complex jobs occupying a safe position (1984, 118-119).

The second important thinker is Caffentzis (2013), who ruminates on the possibilities of capital's tendency towards automation pushed to an extreme. He holds that capital's "dream" of fully-automated "workerless and struggle-less production" entails a "logical escape" for capital, from the annoyances posed by human labour, "through self-reflexivity" or machines which can produce machines without the aid of human labour (2013, 128). Caffentzis holds that: "Only when automata create automata, i.e., when the elements of automata systems become products of automata systems" can high-technology capital "find its fitting foundation" (1993, 128). He notes that while such an "automatization of automation" might seem outlandish, capital continues to make impressive strides today (1993, 129).

The third important thinker is Ramin Ramtin, who attempts to update Marx to the cybernetic era by developing a theory of automation, which he distinguishes from mechanization. Ramtin argues that “‘Automate or die’ is an objective necessity imposed by the very functioning of the capitalist mode of production itself, in accordance with and as a result of the law of value” (1991, 101). Mechanization, which Marx witnessed in his time, is the predecessor to automation, which did not take off until the programmable microprocessor was developed in 1971 (Ramtin 1991, 49). Mechanization involves the machinic copying of human tasks, but automation involves the “systemic application of the principle of feedback” such that a system “includes a procedure of measurement and inspection (or ‘sensing’), the evaluation and processing of this information and an output of instructions as a response, which is then utilized by the system to control all aspects of the particular operations undertaken” (Ramtin 1991, 60). While simple forms of feedback have been utilized since antiquity, the microprocessor enables a new generality of technological recursion. Since it is “a general-purpose logical unit, [the microprocessor] can be programmed to perform an unlimited number of tasks” (Ramtin 1991, 50). Capital’s dream of “self-reflexivity” (Caffentzis 2013, 128) comes one step with closer with the advent of automation because: “there is now a set of technologies available that can actually directly accept objectified information ... without the ... mediation of human labour-power” (Ramtin 1991, 53). Ramtin argues that the most important characteristic of automation is that

it has a capability for communicative and symbolic manipulative simulation ... It is this characteristic which no previous instrument, network of devices or even whole technological systems had ever possessed. It is ... a technology that potentially can be utilized to perform any task which a worker (manual and intellectual) can perform in a production process, and, from capital’s standpoint, much faster and more reliably (1991, 50-51).

Automation presents a possibility that Marx could barely imagine: machines doing intellectual and creative work. Ramtin (1991) notes that “there is increasing pressure to raise the productivity of software specialists ... the main path to the raising of efficiency

here lies in the development of new software packages as a replacement for the skilled labour involved in software production itself” (113). The apex of such software is AI. Ramtin asserts that AI heralds “new wave of automation” (1991, 66) wherein “the process of externalization of control [will become] finally and absolutely complete” (Ramtin 1991, 17). He holds that the “externalization or the complete abstraction of control systems through objectification (in particular as a result of the development of software systems such as ‘expert systems’ or ‘artificial intelligence’) makes possible the removal of the mediatory role of labour. Once this process is fully developed, social and technical control become totally fused” (1991, 58). The increasing sophistication of such technologies “can in fact bring about the objectification of the collective labourer in total. And because of this potentiality, it enables the complete dissociation of living labour ... from the production process” (1991, 58). Capitalism thus tends towards the “ultimate vision of an automated system” (Ramtin 1991, 61): the goal of fusing “conception, coordination and execution into an all-embracing purely *managerial* function” (1991, 65). However, this is not a prediction of capital’s final victory.

On the contrary, Ramtin argues that automation “represents the *final* maturity of the development of the material productive forces under capitalism” (Ramtin 1991, 2, emphasis original) because “automation radically and qualitatively transforms the structure of production, objectively eliminating the process of extraction of surplus labour” (1991, 17). Therefore “the breakdown of capitalism *must* occur with all the inevitable force of a natural law” (1991, 179). Here Ramtin echoes the logic of “The Fragment on Machines”. Next, we turn to another cluster of thinkers who have taken this line of thought up – *operaismo* and *post-operaismo*.

## 2.8 From *operaismo* to *post-operaismo*

Before Braverman popularized the Labour Process approach, a group of theorists and workers in 1950s Italy were applying Marxist ideas to the triad of labour, capital and machine in a different way under the banner of *operaismo* (workerism). Between the 1960s and 2000, the *operaismo* school developed into *autonomismo* (autonomism), and then into *post-operaismo*, a blend of Marxian ideas and poststructuralist theory (see Wright 2002b; Dyer-Witford 1999). Today, *post-operaismo* is widely popular in both

academia and activist circles. It asserts that a qualitative break with orthodox and Western Marxism is necessary due to the computerization and networking of capitalism. Most notoriously, *post-operaismo* has argued that a new form of “immaterial labour” conducted by a new revolutionary subject called the “multitude” means that labour is now technologically-empowered *vis a vis* capital. According to *post-operaismo*, high-tech immaterial labour exhibits a novel autonomy from capital which entails a “crisis of value” (Hardt and Negri 2001). Advocates of this theory thus call for an abandonment of Marx’s theory of value—a position this dissertation is largely devoted to repudiating.

*Operaismo* broke sharply with traditional Marxism and called for more attention to the creative capacities of labour and how they change over time. Thus, *operaismo* theorists did not focus on the deskilling of labour in the face of capitalist machines, but rather on how capital was forced to apply machines in attempts to control the refractory power of labour. The primary theoretical innovation of *operaismo* was to argue for this inversion of the categories of capital and labour. Mario Tronti (1964) held that putting:

capitalist development first, and workers second ... is a mistake ... we have to turn the problem on its head, reverse the polarity, and start again from the beginning: and the beginning is the class struggle of the working class. At the level of socially developed capital, capitalist development becomes subordinated to working class struggles; it follows behind them, and they set the pace to which the political mechanisms of capital’s own reproduction must be tuned. This is not a rhetorical proposition.

Rather than see labour as reacting and resisting capital’s attempts at real subsumption, Tronti and *operaismo* held that it was capital that was reacting to new capacities for insubordination developed by labour. Any new introduction of machines by capital is met with new forms of insubordination such that a series of struggles based on new capacities and with different dynamics can be noted (Holloway 2002, 161-162). *Operaismo* elaborated on the capacities of labour with the notion of “class composition” as a counterpoint to Marx’s notion of the organic composition of capital (Cleaver 1992, 108). Rather than a linear trajectory of ruthless capitalist technological deskilling, *operaismo*

posited a “cycle of struggles” in which capital and labour mutate throughout an ongoing battle, the outcome of which is uncertain (Dyer-Witheford 1999). This perspective aimed to overcome static conceptions of class which posited productive labourers (usually industrial) as the working class and thus revolutionary actor. *Operaismo* instead pointed the way towards a broader conception of the proletariat. Tronti (1966) held that capital “sees society as a means and production as an end”. He and other *operaismo* thinkers used the concept “social factory” (Dyer-Witheford 1999, 134) to describe how society at large was being reconfigured around production.

Antonio Negri pushed this idea further by proposing that a new type of worker was emerging: the socialized worker (Dyer-Witheford 2005, 137-138). The socialized worker is characterized by a world in which society is not only reorganized around the factory, but society as a whole becomes a “factory without walls” (Negri 1989, 89). The locus of antagonism between capital and labour therefore shifts from the factory of the mass worker to society at large. As Dyer-Witheford (2005) puts it, “conflict over exploitation fractally replicates, manifesting in myriad new movements that contest the logic of capital not only in workplaces but also in homes, schools, universities, hospitals, and media” (138). The socialized worker thus has many faces. What is shared by the diverse faces of this emergent figure of labour is an immersion in technology.

Pasquinelli (2015a) shows how early *operaismo* thinkers recognized information technologies as new weapons of capital. Alquati (2013 [1961]) noted how information technologies allowed capital to capture worker’s knowledge and decision-making processes, asserting that the “universal diffusion of capitalist despotism ... realizes itself above all through its technology, its ‘science’”. He was accompanied by Panzieri (2017 [1961]) who saw such technologies as inherently, in their design, antagonistic to labour: “the capitalist use of machinery is not ... a mere distortion of; or deviation from, some ‘objective’ development that is in itself rational”. Panzieri (2017 [1961]) discerned no “occult factor, inherent in the characteristics of technological development or planning in the capitalist society of today, which can guarantee the ‘automatic’ transformation or ‘necessary’ overthrow of existing relations”. There was no self-destruction of capital posited by *operaismo*. The malignancy of technology compelled the young Negri (1979)



to advocate the sabotage of capitalist technology as a necessary condition for worker resistance.

By the era of the socialized worker, Negri's views on machines changed substantially. From advocating sabotage he came to extol technology as empowering labour: "[t]hough initiated for purposes of control and command, as the system grows it becomes an 'ecology of machines'— an everyday ambience of potentials to be tapped and explored by the socialized worker, a technohabitat whose uses can no longer be exclusively dictated by capital" (Negri 1989, 93, cited in Dyer-Witheford 2005, 140). What exactly are these potentials? As Dyer-Witheford (2005) notes, on this Negri was "characteristically abstract" (140), asserting vaguely that productive capacities for the socialized worker center on "science, communication and the communication of knowledge" (Negri 1989, 116). Based on this concept of labour and its new relation to technology, Negri rejected Marx's value theory as a "bourgeois mystification" (Negri 1991, 23). He argued that in high technology capitalism, value is produced socially by the whole of the population, while Marx's theory is based on the work of the individual and is thus "heavily reductive" (Negri 1991, 29) as well as "objectivist, atomized, and fetishist" (Negri 1991, 64).

The notion of the social factory and the primacy of labour over capital, along with Negri's affirmative orientation to technology paved the way for *post-operaismo*, which further elaborated these positions and brought them into widespread use in academia and beyond.

## 2.9 *Post-operaismo*

The most prominent figures in *post-operaismo* are Michael Hardt, Antonio Negri, Paolo Virno, Yann Moulier-Boutang<sup>19</sup> and Maurizio Lazzarato.<sup>20</sup> Lazzarato (1996) laid out the

---

<sup>19</sup> Moulier-Boutang (2012) is an advocate of "cognitive capitalism theory" which differs little from *post-operaismo* and agrees essentially on the theory of immaterial labour. See also Vercellone (2005; 2007)

<sup>20</sup> Other influential *post-operaismo* thinkers include Berardi (2009, 2011), Marazzi (2008, 2011) and Lazzarato (2014). In addition, there is the work of Matteo Pasquinelli – the only *post-operaismo* thinker to

purportedly novel situation of work in post-Fordism and was followed by Hardt and Negri's enormously popular *Empire* (2001). Since then, Hardt and Negri have elaborated their position over the course of three sequels (Hardt and Negri 2005; 2009; 2017). *Post-operaismo* draws not only from Marxist theory but also from poststructuralist philosophers such as Deleuze and Guattari (1983; 1987) and their influences, including Simondon (2011), Freud (1948) and Lacan (1988).

*Post-operaismo* posits a new era of capitalist production called post-Fordism. Post-Fordism is characterized by a new composition of labour called immaterial labour. In theorizing immaterial labour, *post-operaismo* draws on a controversial reading of Marx's "Fragment on Machines". This reliance is of such importance to the theory that Pitts (2017) has suggested calling *post-operaismo* "fragment-thinking" (328). Hardt and Negri (2001) interpret "The Fragment" as a prophecy: "[w]hat Marx saw as the future is our era. This radical transformation of labor power and the incorporation of science, communication, and language into productive force have redefined the entire phenomenology of labor and the entire world horizon of production" (364). But there is one essential difference—capitalism has not imploded. Virno (2001) asserts that "[i]n Postfordism, the tendency described by Marx is actually realised but surprisingly with no revolutionary or even conflictual implication. Rather than a plethora of crises, the disproportion between the role of the knowledge objectified in machines and the decreasing relevance of labour time gave rise to new and stable forms of domination". Despite this, *post-operaismo* remains upbeat. Hardt and Negri (2001) notice in post-Fordism` the "potential for a kind of spontaneous and elementary communism" (294). How is the triad labour, capital, machine configured here?

---

engage directly with AI. He has elaborated the necessity of understanding information machines and data as key components of labour and capital in 21st century capitalism (2009, 2015a, 2017a). He has also contributed philosophical ruminations on AI (2016; 2017b). Despite evidently recognizing that capital and AI have deep, evolving affinities, he has yet to develop a systematic *post-operaismo* theory of AI. Hopefully, this will be provided in his upcoming (2019) monograph.

*Post-operaismo* posits a “change in the quality and nature of labor ... [wherein] information and communication have come to play a foundational role in production processes” (Hardt and Negri 2001, 289). This new kind of labour no longer is subject to Taylorist management. Lazzarato (1996) describes it as “post-Taylorist production” (140). This shift is compelled by the proliferation of the computer, which “proposes itself ... as the universal tool, or rather as the central tool, through which all activities might pass” (Hardt and Negri 2001, 292). The information-based nature of labour enabled by the computer makes labour purportedly immaterial. Lazzarato (1996) defines immaterial labour pithily as “the labor that produces the informational and cultural content of the commodity” (133).

[o]n the one hand ... it refers directly to the changes taking place ... where the skills involved in direct labor are increasingly skills involving cybernetics and computer control (and horizontal and vertical communication). On the other hand ... immaterial labor involves a series of activities that are not normally recognized as ‘work’ – in other words, the kinds of activities involved in defining and fixing cultural and artistic standards, fashions, tastes, consumer norms, and, more strategically, public opinion (Lazzarato 1996, 133).

The first aspect of immaterial labour is digital or informational labour, understood as any work involving the use of computers and/or other ICTs. This is itself a broad category. But it is exceeded in breadth by the second aspect, which Lazzarato sums up as communication (1996, 134-135). He suggests that all of these kinds of work figure into producing the “most important contents” of communication: subjectivities (1996, 140). Taking both of these aspects at once, we might assume that immaterial labour is characterized by an aristocracy of elite high-technology workers who produce information commodities and modulate the subjectivities of the rest of the populace, but this is not the case, or at least it is not only comprised of such workers. Lazzarato holds that “this form of productive activity is not limited only to highly skilled workers; it

refers to a use value of labor power today, and, more generally, to the form of activity of every productive subject within postindustrial society” (1996, 136).

Hardt and Negri define immaterial labour in similarly expansive terms and with a similar dual character. Immaterial labour is sometimes defined as labour “that produces an immaterial good, such as a service, a cultural product, knowledge, or communication” (2001, 290). Hardt and Negri (2001) posit three types of this kind of immaterial labour. The first is industrial production that has been computerized or “has incorporated communications technologies in a way that transforms the production process itself,” such as just-in-time manufacturing (293). The second type is analytical and symbolic work, which may be either “creative and intelligent manipulation” or “routine symbolic tasks”. This refers to the proliferation of computer work, but also to how “[e]ven when direct contact with computers is not involved, the manipulation of symbols and information along the model of computer operation is extremely widespread” (291). The third type is the production and manipulation of affect, which “requires (virtual or actual) human contact, labor in the bodily mode” (293). This includes care work as well as entertainment and any “services of proximity” (293). Most recently, Hardt and Negri restate their position as follows:

the nature and conditions of labor have changed radically from the industrial forms that Marx analyzed ... First, people work in ever more flexible, mobile, and precarious arrangements ... Second, labor is increasingly social and based on cooperation with others, embedded in a world of communicative networks and digital connections ... Capital is valorized through cooperative flows in which language, affects, code, and images are subsumed in the material processes of production (2017, 93).

However, as in Lazzarato’s formulation, for Hardt and Negri, immaterial labour is not limited to a technological elite. At times Hardt and Negri alternatively refer to immaterial labour as “biopolitical labor” or “labor that creates not only material goods but also relationships and ultimately social life itself,” echoing Lazzarato’s notion of the production of subjectivities (Hardt and Negri 2005, 109). *Post-operaismo*’s purportedly

new figure of labour thus has a contradictory dual character. At the same time, immaterial labour refers to both people doing high-tech work with computers and literally every person on earth, as biopolitical producers of subjectivity (Camfield 2007, 47).

On one hand, immaterial labour is performed by the “*social worker*” (Hardt and Negri 2001, 409, emphasis original). The social worker refers to humans throughout society at large, not just the industrial proletariat. The “central role previously occupied by the labor power of mass factory workers in the production of surplus value is today increasingly filled by intellectual, immaterial, and communicative labor power” (Hardt and Negri 2001, 29). These “new productive forces have no place ... because they occupy all places, and they produce and are exploited in this indefinite non-place” (Hardt and Negri 2001, 210). The diverse types of immaterial labour performed by the social worker entails, for *post-operaismo*, a redefinition of what used to be called the proletariat. Hardt and Negri (2001) suggest the term “multitude” which has been widely adopted — but also widely contested.<sup>21</sup> The multitude extends beyond people traditionally classified as workers. It represents the “claim that there is no political priority among the forms of labor: all forms of labor are today socially productive, they in common, and share too a common potential to resist the domination of capital” (Hardt and Negri 2005, 106-107).

On the other hand, immaterial labour is also defined as a specific type of network computer work. Immaterial labour is the result of a “new ‘mass intellectuality’ [having] come into being” (Lazzarato 1996, 134). Moulier-Boutang uses the term cognitive worker or “cognitariat” (96). Thus, for Moulier-Boutang (2012) it is now appropriate to speak of a specifically “cognitive” capitalism which can “only take place on the basis of collective brain activity mobilised in interconnected digital networks” (Moulier-Boutang 2012, 56). Hardt and Negri discern an “artificial becoming” (2001, 216) which produces a “machinic humanity” (2017, 114) for whom “the appropriation of knowledge” and machines “can become decisive” (Hardt and Negri 2017, 119).

---

<sup>21</sup> *Post-operaismo* deploys a term borrowed from Marx (though its meaning is inverted) – the “general intellect” – to describe the technological and social capacities of the multitude (Hardt and Negri 2017, 114). I return to, and critique, the *post-operaismo* uptake of this concept in the final chapter.

Thus, despite the supposed universality of immaterial labour, a particular kind of computerized work is involved in them all: the “computer and communication revolution of production has transformed laboring practices in such a way that they all tend toward the model of information and communication technologies” (Hardt and Negri 2001, 291). The immateriality of computerized work generates a “real homogenization of laboring processes” (2001, 292) in which, for instance, the corporeal care labour of a personal support worker and the code work of a data scientist achieve a new equivalence. This is because capital supposedly no longer cares about concrete labour, but instead is interested solely in “the universal capacity to produce ... abstract social activity and its comprehensive power” (Hardt and Negri 2001, 209). This communicative capacity, in a mutation of Marx’s original definition, as discussed above, is referred to as “abstract labour” (Hardt and Negri 2001, 209). For *post-operaismo*, abstract labour historically follows concrete labour, rather than co-existing simultaneously with it as part of the dual character of labour.

This shows that, as Dyer-Witheford (2005) puts it, “the cyborg, high-tech form of such labor continues to be the privileged point of reference for” immaterial labour theory (152-153). Because of this, I have chosen to evaluate *post-operaismo*’s claims for immaterial labour by exploring the dynamics of work in the AI Industry. As will become clear over the next three chapters, AI work is immaterial labour *par excellence*. Yet, it does not evince certain essential qualities – such as a new autonomy from capital – attributed to it.

Another aspect of *post-operaismo* is worth mentioning here. Because the products of immaterial labour “tend to exceed all quantitative measurement and take common forms,” *post-operaismo* argues that Marx’s labour theory of value can no longer function (Hardt and Negri 2009, 135-136). In other words, because immaterial labour takes on an informational, communicative form, capital encounters a “crisis of measurability” (Marazzi 2008, 43). However, this is not to say that *post-operaismo* abandons the notion of value completely. Hardt and Negri hold that “[e]ven if in postmodern capitalism there is no longer a fixed scale that measures value, value nonetheless is still powerful and ubiquitous” (2001, 356). Value is now produced “in some sense, *externally* to [capital]” (Hardt and Negri 2009 141) by networked “abstract cooperation” which labour conducts

increasingly autonomously from capital (Hardt and Negri 2001, 296). Value thus undergoes a substantial shift in meaning here, along with its associated terminology. In the era of immaterial labour, cooperation “afford[s] labor the possibility of valorizing itself. Brains and bodies still need others to produce value, but the others they need are not necessarily provided by capital and its capacities to orchestrate production” (Hardt and Negri 2001, 294). Valorization, a term Marx uses to describe the process of capital augmentation, is here co-opted by Hardt and Negri to describe a process by which the proletariat increases its own capacities. Value is therefore no longer a term for the average socially necessary labour time for the production of a commodity. *Post-operaismo* therefore calls for a new theory of value which can grasp the “immediately social dimension of exploitation” (Hardt and Negri 2000, 29). I will return to critiques of this, and *post-operaismo*’s immaterial labour theory more generally, in Chapter 6, when I mount my own.

## 2.10 The New Reading of Marx (NRM)

Around the same time as *operaismo* appeared, German-speaking scholars, particularly students of Adorno, were forming a distinct school of Marxism that came to be known as the New Reading of Marx (NRM). NRM is riven by various controversies, but in general, it is characterized by a dismissal of Engels’ work and traditional Marxism as a whole, as “remaining at a purely ‘exoteric’ level that perpetuated traditional paradigms” of political economy (Elbe 2013). Instead, NRM attempts to find the “esoteric” dimension of Marx’s work, which is best described by a focus on the form of value. The emphasis on the form of value leads NRM to assert the ongoing relevance of Marx’s theory of value, contra *post-operaismo*.

Like *post-operaismo*, NRM is a revisionist Marxism, but in contrast to *post-operaismo*’s revisionism, which is founded on perceived empirical changes in the technology of the mode of the production, NRM is based on an exegesis of Marx’s writings enabled by the German language knowledge of early NRM scholars (Pitts 2018b, 1-2; Endnotes 2010). In other words, NRM is based on the claim that Marx has largely been misread, and that in some cases, Marx himself needs to be revised. NRM regards the corpus of Marxian texts, with their many fragments, revisions and version with a “radical openness that

allows [it] to be free of constant reservations based upon what Marx did and did not say on this or that issue, and to move the debate forward into virgin areas of investigation and critique whilst still remaining within in a rich and multifaceted Marxist paradigm” (Pitts 2018b, 24-25).<sup>22</sup>

The primary result of the exegesis conducted by NRM has been the extraction “from the development of Marx’s work a reconstruction of his value theory” (Pitts 2018b, 2) and it as such has sometimes been called value form analysis. Form, in this context, can be understood as referring to a “mode of existence” (Holloway 2002, 51). In the words of Bonefeld, “[v]alue form analysis ... amounts to an exposition of the law of value as a process of social ‘autonomization’, which economics analyses in terms of price movements, stock market developments and other such macro-economic analyses of, in themselves, incomprehensible economic quantities” (2014, 8). NRM mounts its critique not to conduct an economic analysis, but to explain how the system which economics studies is possible – or, how value structures society. But importantly, it does not understand value in an abstract sense: “What is ... ‘autonomized’ is not some abstract essence of value as the ‘ontological foundation of the capitalist system’ that generates an ‘inverted reality’ ... Rather it is the definite social relations of production that subsist in the form of mysterious economic things that seemingly possess the mystic character to ‘instantiate’ themselves” (Bonefeld 2014, 9). Marxism “asks why human social reproduction manifests itself in the form of self-moving economic forces that assert themselves behind the backs of the acting subjects, indifferent and indeed hostile to their needs” (Bonefeld 2014, 21-22). As such, NRM describes itself as the “critique of political economy as critical social theory” (Bonefeld 2014).

NRM’s difference from *post-operaismo* in relation to value is best approached in connection with labour, specifically abstract labour. For *post-operaismo*, abstract labour is the result of a historical homogenization of labour processes wrought by

---

<sup>22</sup> As Pitts (2018b) acknowledges, “Marx’s work on the question of labour and value contains interlaced ambiguities which lend themselves well to varying interpretations, each with its own arsenal of quotations and passages to confirm its position” (24).



computerization. The NRM focuses instead on abstract labour as one of the dual aspects of labour that co-exist in the simultaneous labour/valorization processes that Marx outlined, but with an emphasis on its realization as value in circulation.<sup>23</sup> As Pitts (2018b) puts it, abstract labour is “a category of social mediation expressed in money. It springs from the exchange of commodities by means of money in the sphere of circulation” (3). NRM is not claiming that value is increased only through selling above the cost of making. As Marx outlines, surplus-value is produced via exploitation which occurs in production. However, if the commodity produced never is successfully sold for its socially necessary labour time, then capital never accrues value and the commodity is in effect, without value. For Pitts (2018b) “[a]bstract labour does not so much take place itself, as come about by means of an invention” (26). Heinrich explains this when he states that value “expresses the equal social validity of two completely different concrete acts of labor” (2012, 59). “Prior to being exchanged,” he notes, “the magnitude of value can only be more or less estimated” (Heinrich 2012, 55). Thus, he makes an argument characteristic of NRM: “[v]alue isn’t just ‘there’ after being ‘produced’ someplace” it is a “social relationship ... constituted in production and circulation, so that the ‘either/or’ question is senseless” (Heinrich 2012, 54).

Since neither production nor circulation alone are responsible for value, NRM suggests that *post-operaismo* claims of the collapse of value based on the computerization of production miss the point: “a critical Marxist theory of value situates itself in the circuit of capital as a whole” (Pitts 2018a, 25). In fact, Pitts holds, “postoperaist interpretations of the Fragment’s realization in immaterial labour are seldom immaterial enough” (2018b, 255). Value is, in fact, very much immaterial since it is contingent on realization in exchange. The crisis of measurability posited by *post-operaismo* is based on the notion of value as the expenditure of concrete labour, which when it is socialized across networks of humans by computerization, is rendered impossible to quantify. However, since “value relates to labour’s abstract residue in exchange and not its concrete practice,” this argument fails (Pitts 2018b, 216).

---

<sup>23</sup> Here NRM draws heavily on the pioneering work of Rubin (1972).

*Post-operaismo* thus makes a category error between concrete and abstract labour, or in NRM's terms, between the "'content' of something (its 'natural form')" and its "'social form'" or "'economic form-determination'" (Heinrich 2012, 40). They do not understand value as a "real abstraction resulting from spatiotemporal activity" (Sohn-Rethel 1978, 21, quoted in Toscano 2008, 281). In other words, a real abstraction is an abstraction carried out not in the minds of humans, but one that is "carried out in the actual behaviour of humans, regardless of whether they are aware of it" (Heinrich 2012, 49).<sup>24</sup> One might also approach this notion by saying that the valorization process determines the form of the production process and not the other way around. That is, while *post-operaismo* asserts that changes in the technologies of production entail a disruption of valorization, NRM, on the contrary, asserts that capital valorization adapts technologies to its needs. As Pitts (2018b) puts it, "[c]hanges in the immediate form of labour do not imply changes in forms of abstract social mediation like value" (loc 2391-2401).

NRM does not perceive a radical technological break of the kind *post-operaismo* claims has already happened or will happen. This is because, according to NRM, the fundamental dynamics of the law of value are unaffected by the computerization of labour. Since from a value form perspective, intellectual and manual labour are identical as value, there is no qualitative change when capital begins to rely on scientific and diffuse kinds of labour. Against the purported novelty of immaterial labour, Heinrich holds that "the separation of the intellectual potentialities of the production process from workers is a tendency that is immanent in all capitalist production" (Heinrich 2012, 211). For NRM then, technology thus continues to function as means for relative surplus-value extraction, even with the proliferation of computers and other ICTs. Machines in under capital continue to function, from a value-form perspective, largely the same as they have for the past hundred years or more: "[w]orking through abstraction, capitalism refashions what is real and concrete in the image of the value-form" (Pitts 2018b, 188).

---

<sup>24</sup> NRM is not immune to critique. As Bonefeld (2014) and Pitts (2018b) acknowledge, NRM has often neglected to connect its form analysis to concrete social relations, i.e. class struggle. As Pitts (2018b) emphasizes, it is crucial that NRM analyses recognize that "socially mediate forms are rooted in real relations of antagonism, coercion, domination and dispossession" (5).

This is perhaps why NRM thinkers have devoted little analysis to the particular cutting-edge technologies emerging today, such as AI – an omission which this dissertation intends to rectify. And it is also perhaps why NRM has been silent about *post-operaismo*'s claims for a new technologically-enabled autonomy of immaterial labour, even as the first lucidly devastates the second's claims for a crisis of value. While according to *post-operaismo*'s argument, the new autonomy of immaterial labour is based upon the premise of the crisis of value, this is not necessarily the case. It is possible that labour could experience a new autonomy derived from new technological capacities without the law of value collapsing. For this reason, *post-operaismo*'s claim for a new autonomy of immaterial labour must be investigated in the real world.

## 2.11 Conclusion

In this chapter, I have reviewed several schools of Marxist theory through the conceptual triad: labour, capital, machine. I have done so in order to consider how AI might be analyzed from a Marxist perspective. While AI has of yet received little exploration from Marxist theorists, it has been taken up by a few, but many of these analyses are dated, and there is little consensus to be drawn from them. To arbitrate between these positions, and add to their analysis, it is now necessary to give a fuller account of AI and its relation to capitalism. The next chapter therefore provides a political economic history of how AI came to constitute a powerful sector of capital – the AI Industry.

## Chapter 3

### 3 A Political Economic History of the AI Industry

This chapter charts how AI went from a fringe research interest for a handful of scientists to a feted centerpiece of 21<sup>st</sup> century cybernetic capital. By situating AI in its historical political economic context, we can see it not as an abstract intellectual enterprise or disinterested scientific endeavour, but a technology of automation that has been tightly intertwined with capital, the state and military, since its invention. In the course of this history, I introduce the major schools of AI in their historical social contexts and show how capital has attempted to mobilize them. I show that AI possesses a new level of recursivity, inherent to information technologies generally, which make it an ideal technology for capital, with its inherent drive towards an increasingly machinic form. I discuss the first manifestation of the AI Industry, built around GOFAI expert systems in the 1980s, and how it collapsed due to technological limitations. I conclude by showing how competition between nation states and advances in machine learning AI between the 1980s and 2010s laid the foundation for today's AI Industry, the analysis of which is conducted in Chapter 4.

#### 3.1 The Historical Context

AI is an information technology which runs on computers. If as Ramtin (1991) holds, automation only truly appears with information technology, then any history of AI must also be a history of automation. As Chun (2005) points out, the earliest computers were mechanical; there was no such thing as software and each new program had to be built from physical switches by human hands. The stored program, implemented in the first general purpose digital computer, ENIAC, in 1947, initiated what we today call software, and thus automation (28). The stored program “meant that ‘for the first time it became a practical and attractive proposition to use a computer to assist with the preparation of its own programs’” (Randell, 1973, cited in McCorduck 1979, 62).

Early software was written in binary machine code which was read directly by the computer's CPU. Machine code precisely specifies all functions of a given piece of

software in relation to its hardware. Being composed solely of ones and zeros, machine code was “extremely difficult to use and even more difficult to debug” (Edwards 1996, 247). It was made easier by the development of higher-level programming languages which abstracted from the binary, allowing programming to be done in more-or-less natural languages (Sammet 1969, 1-2). Chun thus refers to programming languages (and programming as we know it today) as the “automation of programming” (Chun 2005, 30) fueled by the “desire ... to recruit the computer into its own operation” (Chun 2005, 29). Software is thus grounded in a process where it is made to work on itself; what Tim Jordan (2015) calls recursion or “the application of information to itself” (31). I suggest viewing the history of AI as the elaboration of this capacity for recursion, and capital’s attempts to appropriate this capacity for its own ends.

The intellectual history of AI is usually traced back at least as far as Alan Turing (see his collected works in Turing 2004). While the influence of Turing is too vast and complex to be adequately explored here, we can briefly note that in 1937, Turing first proposed a hypothetical device now called a Universal Turing Machine. While his point was show that any mathematical problem that could be precisely formulated could be solved by a sufficiently powerful Turing Machine, the paper is also notable for the proposition that human thought could be translated into a series of discrete steps – and thus be subject to mechanical implementation (Edwards 1996, 16).<sup>25</sup> With this gesture, Turing became the philosophical background for all AI.

The story of AI research proper is usually taken to begin less than 10 years after the deployment of ENIAC, in post-World War II Massachusetts, at the 1956 Dartmouth Summer Research Project where the term artificial intelligence was coined.<sup>26</sup> AI

---

<sup>25</sup> Turing is also notable for his famous thought experiment the Imitation Game, proposed in 1948, now often called the Turing Test. In this test, a subject would engage in textual dialogue with a machine and a human, both of which would be hidden from view. The subject would attempt to determine which was which, and if he/she could not, the machine would pass the test. The test is interesting in that it does not strictly test for intelligence, but rather the appearance of intelligence.

<sup>26</sup> Nilsson (2010) points out that two other lesser-known meetings were also important – the 1955 “Session on Learning Machines” in Los Angeles, and the 1958 symposium on the “Mechanization of Thought Processes” at the National Physical Laboratory in Teddington, England.

appeared in a political economic context already enamored with the recursive possibilities of information technology. This was also a context in which the antagonism between capital and labour took on a new, technologically-mediated intensity. While this chapter cannot perform an adequate analysis of the whole situation, four related factors are especially important.

The first is the influence of cybernetics, “the science of control and communication in the animal and machine” (Wiener 1948). Norbert Wiener (1948) popularized cybernetics and its central concern with elaborating connections between the new digital computers and the human brain. Cybernetics drew similarities between humans and machines, based on the use of feedback for control, and invented ways in which the two might be put to work together in complex systems, such as artillery. AI cannot be cleanly separated from the influence of early cybernetics. There was considerable overlap between the two; several of the Dartmouth workshop attendees were involved with cybernetics (Kline 2011, 6). By the 1960s, the fields were distinct, based on a split between “‘symbolic versus continuous systems’ and ‘psychology versus neurophysiology’” (Newell quoted in Kline 2011, 6). While cybernetics focused on continuous systems and physiology, early AI research came to be dominated by the symbolic, psychological approach. Edwards (1996) describes this as a shift from cybernetics’ computer-brain analogy to “the even more comprehensive and abstract computer-mind metaphor of artificial intelligence” – a shift that would not be reversed until at least the 1980s (Edwards 1996, 252). The precise interactions between AI and cybernetics are too complex to chart here, but we can note that cybernetics, by breaking down conceptual barriers between human and machine, contributed to the creation of a “cyborg discourse” (Edwards 1996, 2). AI would take up, and advance this discourse in its quest for increasing recursivity.

A second factor is the increased industrial productive capacity of capital driven by advances in automated machinery and computing technology. As early as 1946 *Fortune* magazine published an influential article entitled “Machines without Men” which called for the full-scale application of wartime technological advances in automatic machinery to peacetime manufacturing (Leaver and Brown 1946). Two years later, discussing the industrial endeavours of the Ford Motor Company, Le Grand (1948) produced the first definition of automation as “the art of applying mechanical devices to manipulate work

pieces into and out of equipment, turn parts between operations, remove scrape and to perform these tasks in timed sequence with the production equipment so that the line can be wholly or partially under push-button control at strategic stations”. The exigencies of military production in World War 2 had driven a massive increase in productive capacity in the Allied countries, particularly the USA and Canada. Increased capacity was due to a growth in the sheer number of factories but also to the intensification of labour processes via Fordism and the wide application of mechanization and automation technologies. According to Smith (2000), the “height of the Fordist era coincided with the first period of the so-called computer age. Beginning in the 1950s corporations introduced mainframe computers for data processing” (5). Computers enabled the introduction of information technology automation such as numerical control machine tools (Noble 2011). The first computerized automation technologies were installed in chemical plants in 1960 and became prevalent over the course of the decade. In 1959, the first industrial robot, called Unimate, was installed at a General Motors factory in New Jersey. It was a relatively simple machine by today’s standards, consisting of a single arm which could perform one task at a time. It was in mass production by 1961 and was widely copied, leading to industrial robot proliferation (Wallén 2008, 9-10). Machines performing actions previously reserved for humans in both manual and (limited) cognitive contexts were now a reality and their economic potential was beginning to be realized.

The third factor is the changed face of post-war labour. The mass worker of the Fordist workplace came out of the Second World War in a strong position *vis a vis* capital. The USA’s no-strike pledge had expired and “newly powerful unions [were] ready to test their strength” against capital (Noble 2011, 24). However, labour’s power, which manifested in numerous, often wildcat, strikes over wages, drove capital to conduct a multipronged assault on labour, one powerful arm of which constituted the application of automation technologies. As Noble (2011) puts it, the post-World War II years offered capital “unprecedented opportunities” for “reduced skill requirements, more concentrated management control, and the replacement of workers by machines” (36).<sup>27</sup> While, as

---

<sup>27</sup> Norbert Wiener himself recognized the long-term potential effects of automation on workers. In an unanswered 1949 letter he wrote to then UAW president Walter Reuther, Wiener warned that progress in

Noble (2011) demonstrates, automation technologies did not achieve the delirious goals of perfection their promoters promised, they did allow firms to reduce the number of unskilled and low-skilled workers they employed. And as *operaismo* theorists have emphasized, this had the effect of breaking the power of unified mass workers (Wright 2002a). At the same time, many capitalist states turned to Keynesian welfare to effect a “class compromise” between labour and capital (Harvey 2007, 10). This successfully blunted the antagonistic energies of post-war labour. Industrial workers traded in revolutionary aggression for benefits and wage increases.

Outside of the factory, other changes to labour were occurring. The management theorist Peter Drucker (1957) pronounced the rise of a revolutionary new knowledge society founded on the automation of manual labour and the technologically augmented skills of the knowledge worker. One year after the Dartmouth workshop, Drucker argued that:

Today the assembly line is obsolescent ... even mechanical work is best organized as joint effort of men of high skill and knowledge exercising responsible, decision-making, individual judgment in a common effort and for a joint end ... Automation may well eliminate the unskilled worker from the production floor. But it replaces him by an equal number of men of high skill and judgment ... Each of them works in his own field of knowledge with a broad discretionary area of judgment. Each of them, however, must of necessity work closely with all the others - in constant communication with them, constantly adjusting to their decisions and in turn making decisions that affect their work (Drucker 1957, 67).

---

automation “will undoubtedly lead to the factory without employees”. He goes on to state: “I do not wish to contribute in any way to selling labor down the river, and I am quite aware that any labor, which is in competition with slave labor, whether the slaves are human or mechanical, must accept the conditions of work of slave labor” (Wiener 1949)



Looking back, we now know that what Drucker observed applied only to a small subset of workers in only a few parts of the world. High-skill collaborative and communicative work did appear, but it did not do away with previous forms of deskilled work. Yet, discourse such as Drucker's continues to appear alongside technological innovations, as we will see. Indeed, even at the time of Drucker's writing, new fields of work were appearing which exhibited none of the shining characteristics he describes.

As the mass worker squared off with increasingly powerful automated industrial machinery, a new workforce was forming around emerging computing technology, and was being, even at that early stage, subject to deskilling forces analogous to those experienced by industrial labour. Computer scientists and engineers were few and a large amount of routine programming and debugging labour was required to make computers actually work. By 1955, "the one thousand or so extant general-purpose computers required the services of perhaps 10,000 programmers. Five years later, in the midst of a booming commercial computer market, programming had suddenly become a profession in its own right, with about 60,000 practitioners servicing some five thousand machines. Programming began to emerge as a craft" (Edwards 1996, 248).

However, even as programming was emerging as a craft, "already the amount of mathematical skill it required had begun to diminish" (Edwards 1996, 248-249). In its early days during the Second World War, computer work was "the job of the dispossessed, the opportunity granted to those who lacked the financial or societal standing to pursue a scientific career. Women probably constituted the largest number of computers, but they were joined by African Americans, Jews, the Irish, the handicapped, and the merely poor" (Grier 2005, 276). Ten years after the end of war, this work was still performed by subordinated groups: "[m]usic teachers and women without specialized backgrounds were among the most successful [programmers]. Such groups could not be expected to learn machine code or produce mathematically elegant algorithms; to make this new work force effective required symbolic languages easily learned by nonspecialists" (Edwards 1996, 249). Deskilling was further motivated by the fact that "[b]usinesses ... wanted to write their own software without hiring expensive experts ... they had to if they were going to use computers at all, since before the 1960s 'off-the-

shelf' software was virtually unknown" (Edwards 1996, 249). The emerging programming workforce was stratified from its inception. I will show later that this quality has continued into the contemporary AI Industry.

While far from a complete picture, these factors outline the context of technologically-reconfigured antagonism between labour and capital in which AI emerged - and from which its critical study should not be divorced.

### 3.2 The Advent of AI

The organizers of the 1956 Dartmouth Workshop (funded by the Rockefeller Foundation and the US Office of Naval Research (ONR)) founded their project:

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, Minsky, Rochester and Shannon 1955).

These AI pioneers set out to achieve the same capture and emulation of human activity that Marx posited as the function of automatic machinery, but in the mental, rather than physiological, domain. It is therefore not hard to imagine why capital might find AI appealing. Only a few years later, two influential researchers proclaimed that "there are now in the world machines that think, that learn, and that create ... their ability to do these things is going to increase rapidly until in a visible future-the range of problems they can handle will be coextensive with the range to which the human mind has been applied" (Simon and Newell 1958, 8). They go on to explain that "[w]ith recent developments in our understanding of heuristic processes and their simulation by digital computers, the way is open to deal scientifically with ill-structured problems - to make the computer coextensive with the human mind" (Simon and Newell 1958, 9). The Dartmouth organizers even hoped for a recursive kind of self-modifying AI, noting that "[s]ome schemes for [self improvement] have been proposed and are worth further study".

They were terribly optimistic, holding that their “2 month, 10 man study” could make a “significant advance” towards machine intelligence (McCarthy, Minsky, Rochester & Shannon 1955).

Such lofty aspirations proved exorbitant; few of their goals were realized. Hopes for cybernetic learning machines inspired by neurophysiology were largely dashed, beyond some interesting curiosities, and instead, an approach now known as GOFAI (Good Ol’ Fashioned AI) (also known as symbolic AI) took off and remained, until the mid-1980s, “the dominant (though not only) approach in AI” (Boden 2014, 89). GOFAI is an approach to AI that aims to implement high-level cognitive functions, such as logical reasoning, in machines through the manipulation of information encoded in a symbolic language. In contrast to learning machines, GOFAI “sought first to formalize knowledge of the world, injecting it into computer systems predefined and predigested. Logic, not experience, would determine its conclusions” (Edwards 1996, 255).

GOFAI is thus based on the assumption that “a large part of human thought consists of manipulating words according to rules of reasoning and rules of conjecture” (McCarthy et al 1955). It can also be summarized by the assertion that “[s]ymbols lie at the root of intelligent action” or that intelligent action is produced by the manipulation of symbols (Newell and Simon 1976, 114). A GOFAI system solves problems by “generating and progressively modifying symbol structures until it produces a solution structure” (Newell and Simon 1976, 120-121). This is also called “heuristic search” (Boden 2014, 90). This means that the system creates internal representations of its world or problem domain in a symbolic language and performs logical or rule-based manipulations on this representation to solve a problem.

One obvious example here is chess. There are a finite (though large) number of possible states the chess board can be in and a finite set of rules that determine its possible state changes and future states. There are a finite number of possible moves at any given time. However, a good chess player looks not only at the current state of the board, but also considers possible future states, and plans ahead. However, mentally scrolling through all

possible moves both current and future is time-consuming and cognitively wasteful, so a good chess player will apply some form of heuristics, or rules of thumb that may not always work properly, but cut down the list of options to be searched. A chess-playing GOFAI system will also employ heuristics.

The necessity of explicit internal representation led to a serious issue for GOFAI, known as the frame problem, or “the problem of updating, searching, and otherwise manipulating, a large structure of symbols in realistic amounts of time” (Copeland 2000, see also Dennett 2006). One way to think of this is in terms of common sense reasoning. For instance, if we humans perceive a plate resting on a table (in the absence of contributing factors such as a meddlesome cat) we do not wonder whether it might suddenly disappear, or reappear under the table, or suddenly shatter. We know that force must be applied to it for any of these things to happen. A GOFAI system must have all such basic notions programmed into it explicitly in the form of rules, and there are very many rules required to accurately approximate even naïve conceptions of physics. Since all of these rules had to be discerned and then programmed, beyond the simplest virtual domains, GOFAI was bound to run into problems of complexity.

Despite this and other problems, most of AI’s early successes can be credited to GOFAI and it was GOFAI that first constituted the AI industry, in the form of expert systems, as I will show. While GOFAI’s fame has since faded, it is still used in the AI subfield of planning which finds applications in video games, route planning, air traffic control, and configuring electronics components (Kaplan 2016, 25). Few now see GOFAI as the future of AI since, as Boden reflects, “most of the ‘intelligence’ involved [in GOFAI] lies in the choices of actions, operators, and heuristics specified by the programmer” (2014, 90). But in the early years things were different.

There was not yet an AI industry (ie. AI was not yet sold as a commodity) in the 1950s and 60s, although those researching the topic came from both industry and academia – usually, with the majority of funding from the US military. By the end of the Second World War, the US Office of Scientific Research and Development (OSRD) had increased military research and development funding from a prewar ~\$23 million to over \$100 million yearly (Edwards 1996, 46). This generous outlay of research funding

continued after the war with the establishment of a small number of “centers of excellence” in research, situated in universities (Edwards 1996, 261). Some of these were devoted to AI. Starting in the 1960s, “AI, for over two decades almost exclusively a pure research area of no immediate commercial interest, received as much as 80 percent of its total annual funding from ARPA” (Edwards 1996, 64).

The Advanced Research Projects Agency (ARPA) and the Office of Naval Research (ONR) were two of the largest funders of AI research. Formed in 1946, the ONR is a department of the US Navy that engages in and funds technology research and development with applications to naval security. The ONR’s Information Systems branch funded several AI projects at a number of institutions (Nilsson 2010, 118). ARPA was formed in 1958 as a response to the USSR’s successful deployment of the Sputnik 1 satellite the year prior. The agency changed its name to Defense Advanced Research Projects Agency (DARPA) in 1972 and 1996 and reverted briefly to ARPA from 1993-1996. The agency’s overall task remains the achievement and maintenance of US technological superiority through organizing and funding research; including AI: “DARPA has held to a singular and enduring mission: to make pivotal investments in breakthrough technologies for national security” (DARPA 2015, 1). ARPA opened an Information Processing Techniques Office (IPTO) in 1962. Directed by J.C.R. Licklider, who advocated a “symbiosis” between humans and machines (Licklider 1960), the IPTO funded AI and other computing projects at MIT, Stanford, Carnegie Mellon, SRI, RAND, BBN and SDC, among others, beginning around 1962 (Nilsson 2010, 120).

It is difficult to categorize the researchers from this period as belonging to either academia or industry. Most of them passed between the two throughout their careers, often more than once. The integration of academia and industry, built on the foundation of military funding leads Edwards (1996) to categorize the nascent computing industry as an “‘iron triangle’ of self-perpetuating academic, industrial, and military collaboration” (47). This dynamic continues today in the AI Industry.

The universities associated with early AI research are few: MIT in Cambridge, Massachusetts, Carnegie Tech (now Carnegie Mellon University) near Pittsburgh,

Pennsylvania, Stanford University in Stanford, California and the University of Edinburgh in Edinburgh, Scotland.

The industry organizations most important to AI's inception were all based in the USA. Perhaps the most well-known of these is the RAND (Research and Development) Corporation headquartered in Santa Monica, California. RAND still exists as a non-profit organization which conducts science, technology and policy research. Originally Project RAND, a 1945 US Air Force project with a mandate to direct research and development towards strategic ends. The project was spun-off into its independent form in 1948 and developed into one of the most influential private institutions of the past century. As one RAND researcher puts it in a popular history of the organization, “[s]atellites, systems analysis, computing, the internet – almost all the defining features of the information age were shaped in part by the RAND Corporation” (Campbell 2004, 50). The first book on AI, the anthology *Computers and Thought* (Feigenbaum and Feldman 1963) had no less than six of its twenty chapters previously published by RAND researchers (Klahr and Waterman 1986, 1).

International Business Machines (IBM) has also been involved with the development of nearly all aspects of information technology, AI included. Founded in 1911 as “Computing-Tabulating-Recording Company” in New York City, it was renamed IBM in 1924 and moved its headquarters to Armonk, New York in 1963. Arthur Samuel, an IBM researcher, developed a checkers playing program which was able to improve itself and successfully won a televised match in 1956.<sup>28</sup> Samuel is attributed with the first use of the term machine learning. He hoped that “[p]rogramming computers to learn from experience should eventually eliminate the need for much of this detailed programming effort” (Samuel 1959, 211).

Bell Laboratories was also involved in the development of a myriad of information technologies, but also lasers and solar panels, as well as early experiments in AI. Claude Shannon, inventor of information theory, worked for Bell starting in 1941 and developed

---

<sup>28</sup> While Samuel is often cited as the creating the first checkers playing program, Nilsson (2010) notes that Christopher Strachey at the University of Oxford seems to have created one as early as 1951 (90).

an artificially-intelligent mouse called Theseus based on technology developed for telephone switching. The mouse employed a rudimentary form of machine learning to find its way out of mazes.

The Systems Development Corporation (SDC) is considered the world's first software company (Campbell-Kelly 2004, 36-41). It began as a project group of the RAND Corporation in 1955 and was spun off as a non-profit in 1957. SDC specialized in designing and consulting on large and complex, networked computer systems. They are perhaps best known for their involvement in the production of the Semi-Automated Ground Environment (SAGE). SAGE was a system of networked radar sites and computer that was intended to provide comprehensive surveillance of US airspace for the purposes of defense. Because the amount of the data the system was required to process was enormous, AI techniques were deployed. SAGE went live on 1958 and was never put to the test. In 1969, SDC turned into a for-profit company.

A number of other firms were influential. BBN Technologies, founded in 1948, in Cambridge, Massachusetts, was led to AI by the complex computations involved in acoustic design research. BBN was involved in the development of several influential technologies including the internet. SRI International (formerly Stanford Research Institute) is a nonprofit thinktank established by trustees of Stanford University in 1946. In 1966, the institute opened its Artificial Intelligence Center that developed "Shakey," widely accepted as the first mobile robot with the ability to perceive and reason about its surroundings (SRI International n.d.). Lincoln Laboratories, created in 1951 as a federally-funded research and development center at MIT, was meant to remedy perceived inadequacies in the USA's air force. The Laboratories' famous first project was the SAGE system which SDC also contributed to.

Early AI researchers moved between these institutions throughout the 1950s and 60s. There seems to have been no discernable trend in migrations from academic to industry positions or vice-versa. J.C.R. Licklider, for instance, worked first at Lincoln Laboratory and MIT and later went to BBN and eventually, to ARPA (discussed below) (Nilsson 2010, 119). Allan Newell on the other hand, worked at RAND in his early years but spent his later years solely at CMU (Nilsson 2010, 115). While there were, by the late 1950s

and early 1960s, labs devoted to AI at some of the above universities, researchers would often have to use the powerful and expensive computers owned by companies such as BBN for computationally demanding AI projects (Nilsson 2010, 116). MIT owned an IBM 704 which was cutting-edge at its release in 1954, but it quickly began to pale as AI research intensified. The first PDP-1 (programmed-data-processor) computer which possessed magnitudes more computational power was bought by BBN in 1960 for USD\$120,000 (USD\$1,041,100.34 in 2019 dollars).<sup>29</sup> This expensive hardware ensured that AI researchers would continue to be reliant on the deep pockets of military funding.

This situation obtained until the late 1960s. Military-funded AI research flourished in the limited context of think-tanks and academic departments. Much of the work done at this time was based in the GOFAI paradigm, but some research was also conducted into rudimentary artificial neural networks, which I return to later. There were a number of impressive achievements in fields such as machine vision, robotics and theorem proving. There were, however, no commercial applications developed.

### 3.3 The AI Winter

The end of the 1960s and early 1970s were a dark time for AI research. While the world economy began tapering off from the post-WWII boom, “stagflation” threatened Western states and civil rights activists mobilized for a variety of causes including gender relations and opposition to militarism, the excessive optimism of proclamations about AI, such as those made by the Dartmouth researchers, became recognized as such. A barren “winter” ensued in which funding for AI research dried up (Crevier 1993, 203). In November 1966, AI research received a hefty blow by the negative assessment of the ALPAC (Automatic Language Processing Advisory Committee) Report, commissioned by the US government to evaluate the prospects for machine translation research. The report ended most government funding into machine translation and damped enthusiasm for AI in the public sphere. In 1969, the US Congress passed, as part of the 1970 Military Authorization Act, the first Mansfield Amendment. Section 203 of the amendment stated:

---

<sup>29</sup> Adjustment for inflation according to US Bureau of Labor Statistics:  
[https://www.bls.gov/data/inflation\\_calculator.htm](https://www.bls.gov/data/inflation_calculator.htm)



“None of the funds authorized by this Act may be used to carry out any research project or study unless such project or study has a direct and apparent relationship to a specific military function” (quoted in Laitinen 1970, 689). The Mansfield Amendment had the effect of strictly limiting the kinds of AI research government bodies, such as the ONR, could fund. The amendment “brought about a swift decline in some of the military’s support for basic research, often driving it toward the applied realm” (National Research Council 1999, 213). ARPA’s autonomous nature made it exempt for the 1969 amendment, but a second amendment, passed in 1973, applied the same restrictions specifically to that agency. Despite a restricted range of research problems, by the mid-1970s “as much as 80 to 90 percent of funding for the major AI groups – MIT, Stanford, SRI and Carnegie-Mellon” was provided by ARPA (Edwards 1996, 296)

1973 also saw the publication of the Lighthill Report in the UK. The mathematician Sir James Lighthill was commissioned by British Science Research Council to study the prospects of AI. The report, “Artificial Intelligence: A General Survey” was damning and recommended that much AI research was pointless as it would remain stymied by the fundamental problem of the combinatorial explosion. This report dampened AI research in the UK, which did not recover until the early 1980s, and arguably contributed to American dominance in the field.

Despite the climate of the AI winter, some universities opened new AI research groups. The most prominent were Toronto, Rochester, Texas, Maryland, British Columbia, California and Washington, but there were others in Europe and Asia (Nilsson 2010, 207). But with the major sources of funding for AI research now strictly limited, many academic researchers working on basic research were compelled to turn to applied research while others left academia altogether. This chilling effect is often credited with contributing to the rise of startups, the proliferation of private research labs and the computing industry in general. Non-AI computing companies, such as Microsoft, founded in 1974, flourished at this time. Microsoft released the influential operating system MS-DOS in 1981. Victor (2013) suggests that “[o]ne way of interpreting ... the Mansfield Amendment [is that it] killed research, but ‘induced labor’ on an industry”.

### 3.4 The Birth of the AI Industry: Expert Systems

*To make major productivity gains, we discovered that the automation net would have to be cast farther out than the factory floor – to cover the information handlers, those doing the planning, the problem solving, and the decision making. In short, it was necessary to bring the power of automation to the knowledge worker (Feigenbaum, McCorduck and Nii 1988, 4)*

In the 1970s, social theorists such as Daniel Bell took up and elaborated ideas broached first in Drucker's notion of the knowledge worker. Bell (1973) predicted a coming "post-industrial" society. This would be characterized by "a shift from a goods producing to a service economy; a move in occupational distribution away from manual labour to the pre-eminence of professional and technical work; increasing capacities of assessment and forecasting; and a new 'intellectual technology' of games-theory and systems-analysis, materially embedded in computer systems" (Dyer-Witheford 1999, 30).<sup>30</sup> Despite the tendency of epochal modes of analysis like Bell's to overemphasize change at the expense of continuity (Webster 2014), there did occur substantial changes around that time. Between 1948 and 1990 employment in the USA increased by nearly 83% with 97% of this increase coming from "nongoods-producing industries" including "services, transportation, communications, utilities, wholesale and retail trade, finance, insurance, and real estate" (US Bureau of Labour Statistics 1993, 7). The shift towards service jobs truly took off only in the 1970s. Between the 1970s and 1990s, "the process of economic restructuring and the technological transformation ... led to a reduction in the share of manufacturing employment in all countries" which varied considerably, but in some countries surpassed 20% (Castells and Aoyama 1994, 11).

---

<sup>30</sup> Marx-influenced critical analyses would use the notion of post-Fordism to discuss the very same phenomena, as the previous chapter discussed (see also Amin 1994). *Operaiismo* theorists such as Negri (1989) proposed that the mass worker of Fordism was being replaced by a new figure called the socialized worker. The socialized worker referred to diverse people living in a world "where capital had insinuated itself everywhere," and work now depended more on communication and sociality than it had before (Dyer-Witheford 2005, 138). By the publication of *Empire*, the socialized worker was reconfigured into the "immaterial" labourer (Hardt and Negri 2001, 290) who was adept with information technology and might even be able to turn it against capital.

Coincidentally or not, as skilled, intellectual labour burgeoned, so did efforts to automate it. As capital reconfigured itself around so-called knowledge work during the 1970s, AI, for the first time, coalesced into an industry based on an application of GOFAI called expert systems or knowledge-based systems. Expert systems were intended to capture sophisticated knowledge about a particular expert domain and make it available to management and/or workers lacking in said knowledge. While some of these systems attempted to embody the knowledge represented by textbooks or documents, most often, they were intended to model the knowledge of expert human workers and thus explicitly manifested the Taylorist principle of “the dissociation of the labor process from the skills of workers” (Braverman 1998 [1974], 78). Critical analysts of the day did not fail to notice this, as I show below.

Expert systems are comprised of two basic parts: the knowledge base and the inference engine. The knowledge base is comprised of factual and heuristic knowledge, represented in a symbolic language. It is created by a process called knowledge engineering in which human engineers interview experts in a given field to gather their knowledge and then encode this knowledge in a symbolic language as conditional (if-then-else) rules (Feigenbaum, McCorduck and Nii 1988). The inference engine defines the steps taken to apply the stored knowledge. The idea is that once the knowledge is encoded into the expert system, less skilled workers or management can simply solve problems in that domain, without recourse to the former possessor of that knowledge.

Intellicorp, arguably the first AI startup, was founded in 1980 by Edward Feigenbaum and others. Intellicorp “won a contract in 1984 for \$1,286,781 [from DARPA] to develop an ‘evolutionary new generation system tool’” (Roland and Shiman 2002, 198). In 1984, DARPA “budgeted \$1,813,260 for Teknowledge over two years, with an option to renew” (Roland and Shiman 2002, 201). Teknowledge was another startup founded in 1981 by Feigenbaum. Along with some nineteen scientists from Stanford, as well as Frederick Hayes-Roth, then Research Program Director for Information Processing Systems at RAND, Teknowledge began as a consulting business, but with ARPA’s funding moved into expert systems. This interplay of academia and industry, boosted by military funding, brought the AI industry into being and continues to characterize the AI industry today.

The first expert system, Dendral, was developed throughout the 1960s by Feigenbaum and Joshua Lederberg at Stanford University with the goal of helping organic chemists identify unknown biological molecules (Lederberg 1987, 1). While it mainly just eliminated implausible compounds, Dendral proved that the expert system concept could work, and while it remained in academia, it spawned many offspring, some of which were commercialized. The earliest of Dendral's descendants to be deployed commercially were developed in cooperation between academia and industry. The "Dipmeter Adviser" was developed by the oilfield services corporation Schlumberger in cooperation with academic AI engineers. This system was used for "inferring subsurface geologic structure" in preparation for drilling by analyzing data from sensors lowered deep into the earth's crust (Davis et al., 1981, 846). A prototype was developed in 1980 but the system was not deployed in Schlumberger's operations until 1984 (Davis et al., 1981, 129). The system had 90 rules which used about 30 predicates and functions to represent data (Smith and Young 1984, 16). In 1980, Digital Equipment Corporation (DEC) deployed the "first expert system in daily production use in industry," XCON (Expert Configurator) (Virginia and Dennis 1989, 298). XCON aided in the configuration and assembly of complex, customizable computer systems (Virginia and Dennis 1989, 221). Compared to Dipmeter Adviser, it was massive, containing over 2000 rules (Kraft 1984, 43). According to Roland and Shiman (2002), by "the middle of the decade [DEC] estimated it was saving \$40 million annually by use of XCON" (191). Word of the utility of expert systems spread throughout the business world and demand exploded throughout the mid-1980s (Crevier 1993, 198). Commenting on the "industrialization of artificial intelligence" in 1984, Kaplan notes "[o]ver the past few years, the character of the AI community has changed. AI researchers used to be able to go about their work in peace ... As the promise of practical applications of AI has slowly become reality, new players have entered the field, changing its nature forever" (1984, 51). According to one contemporary estimate, 1500 expert systems were used commercially by the end of 1987 (Feigenbaum, McCorduck and Nii 1988, x).

It is only at this point in the late 1980s that it is fair to say that AI comprised a distinct industry. Crevier describes it as composed of three main sectors. The largest sector, comprising over half the market, were specialized AI computers called LISP machines,

the largest producers of which were the corporations Xerox and Texas Instruments and two startups from Massachusetts: Symbolics and Lisp Machines (Crevier 1993, 200). The second largest sector was that of expert system “shells” or frameworks which had to be engineered to function on its particular knowledge base. The third sector, complete expert system applications, were the smallest sector of the industry. Roland and Shiman (2002) note that the “commercial firms springing up in the early 1980s were building custom systems one client at a time. DARPA would try to raise the field above that level, up to the generic or universal application” (Roland and Shiman 2002, 194). Expert systems were very narrow, but DARPA wanted a generally-applicable expert system that could be deployed to different domains without its core components needing to be modified. The development of such a generic expert system guided research programs at both Teknowledge and Intellicorp. While such a recursive expert system was not achieved, interest in, and funding for, expert systems drew labour to the field. The number of people trained to work in AI dramatically increased throughout the 1980s due to the establishment of new graduate programs aimed not at producing academic researchers, but instead nurturing the basic skills for practical applications of AI. Feigenbaum developed such an applied Masters’ of Science in Artificial Intelligence that was the first of its kind (Roland and Shiman 2002, 196). Membership in the Association for the Advancement of Artificial Intelligence (AAAI) “rose from around 5,000 shortly after the society’s founding [in 1979] to a peak of 16,421 in 1987” (Nilsson 2010, 271). The increase in skilled labour enabled “[m]ost large companies [to establish] AI groups to develop in-house applications. In 1985, 150 companies spend [sic] \$1 billion altogether on internal AI groups. At DEC, for example, the AI group had grown to 77 people by 1986, and mushroomed to 700 in 1988” (Crevier 1993, 199). In 1987, the conglomerate DuPont “had 100 expert systems in routine operation and 500 in various stages of development” (Crevier 1993, 199).

While detailed figures are hard to come by, Crevier (1993) gives one snapshot of the nascent AI industry. According to Crevier, in 1986, sales in the USA of “AI-related hardware and software reached \$425 million (1993, 200). Forty new companies were formed that year alone, with total investments of around \$300 million (Crevier 1993, 200). Enthusiasm for expert systems continued throughout the 1990s. From 1500 expert

systems in use in 1987, the AI industry exploded to an estimated 12,500 expert systems in 1997 (Liebowitz 1997, 118).

### 3.5 Strategic Computing: AI and the State

The 1980s have been characterized politically predominantly by the rise of neoliberalism, first in Chile, the UK and USA, and then elsewhere. Harvey defines neoliberalism as a “theory of political economic practices that proposes that human well-being can best be advanced by liberating individual entrepreneurial freedoms and skills within an institutional framework characterized by strong private property rights, free markets, and free trade” (Harvey 2007, 2). While neoliberal theory extols the virtues of market freedom from state influence, it has in reality entailed the formation of a “neoliberal state” (Harvey 2007, 7) which implements whatever practices are necessary to keep the market operating. Neoliberal states have shown an “intense interest in and pursuit of information technologies” because they require “technologies of information creation and capacities to accumulate, store, transfer, analyse, and use massive databases to guide decisions in the global marketplace” (Harvey 2007, 3). One manifestation of this was state involvement in AI research in the 1980s, carrying initial military interest into commercial domains. While neoliberalism was spread unevenly then, its technological impetus was contagious. The world’s most powerful states then engaged in a pursuit of AI technology not unlike the Space Race, even if it ended much less spectacularly. AI did not become an industry in its own right solely because a few intrepid academics, such as Feigenbaum, ventured out into the business world. During the 1980s, AI (primarily expert systems) and computing became axes of competition for the most technologically-sophisticated states.

Japan’s Ministry of Trade and Industry (MITI) jumpstarted the process with its Fifth Generation Computer Systems project in 1982. MITI created the Institute for Next Generation Computer Technology (ICOT) to carry out a ten-year plan of developing massively parallel expert systems specifically for implementing AI (Feigenbaum and McCorduck 1984, 31-32). ICOT was backed by a consortium comprised of two government research labs and eight Japanese giants of industry: Fujitsu, Hitachi and Toshiba (Feigenbaum and McCorduck 1984, 130). MITI committed to \$450 million over

the ten years with large contributions from the firms composing a budget of approximately \$850 million (Feigenbaum and McCorduck 1984, 137). By 1992, some interesting technological developments had been achieved, but commercial success had not. While the project did not achieve its goals, it did substantially advance Japan's computing infrastructure and knowledge, though not enough to contest the dominance of the USA (Roland and Shiman 2002, 326).

Reacting to Japan's computational ambitions, the USA began its Strategic Computing (SC) Initiative in 1983.<sup>31</sup> By 1993, DARPA had spent "an extra \$1 billion on computer research to achieve machine intelligence" (Roland and Shiman 2002, 1). The main goal was generic expert systems employing massively parallel processing: "AI would become an essential component of SC; expert systems would be the centerpiece" (Roland and Shiman 2002, 192). It is under the aegis of SC that Teknowledge and Intellicorp received the funding they did, as discussed above. SC also failed to meet its ambitious goals for AI, and was instead adapted to a project of "high-performance computing" by the end of the 1980s (Roland and Shiman 2002, 325). Like Japan's Fifth Generation project, it did however contribute greatly to improved computing infrastructure and a proliferation of knowledge about AI (Roland and Shiman 2002, 331). Not all USA AI initiatives failed, however, the Dynamic Analysis and Replanning Tool (DART) was developed for DARPA by a team led by an engineer from BBN. Deployed in 1991, it successful logistical management reportedly "paid back all of DARPA's 30 years of investment in AI in a matter of a few months" according to the director of DARPA at the time (Hedberg 2002, 81-83).

The USA's technological mobilization was also motivated by fears of the increasingly technological figure of communism. There is little available literature in English about Soviet AI research, although it is certain that some did occur. Computing got off to a slow start in the USSR due to at least two factors (Peters 2016). One is that the Party

---

<sup>31</sup> 1983 also saw the advent of the ballistic missile oriented Strategic Defense Initiative (SDI). The connections between the two initiatives were subject to much debate although the SDI was explicitly described as responding to the Soviet, rather than Japanese, threat (Roland and Shiman 2002, 63-64).

concealed the existence of computers in the Union from both its citizens and other nations for the first years of their existence, slowing computational skill diffusion. Second is that cybernetics was construed, well into the mid-1950s, as an ideological taboo for Soviet cold warriors, a bourgeois pseudo-science that contradicted Marxism-Leninism (Gerovitch 2004). After the death of Stalin, cybernetics gradually became acceptable and research into AI became acceptable. According to Malinovsky (2010), the prominent Soviet cyberneticist Glushkov “regarded artificial intelligence as the most promising direction in Cybernetics” (32). Soviet AI practitioners pursued research into theorem proving, pattern recognition and machine vision (Gerovitch 2011, 180; Peters 2016, 117). In 1960, a visiting RAND researcher surveyed the Soviet AI and computing landscape and noted several AI projects underway including learning machines, speech recognition, heuristic search and machine translation. He concluded that in the USSR “computer time is becoming less scarce and that research on artificial intelligence and other advanced applications of computers is beginning to achieve a priority which it did not seem to have just one year ago” (Feigenbaum 1961, 579). By 1966, at the first machine versus machine chess championship, a Soviet chess-playing AI program defeated by 3-1 a program built by AI practitioners from MIT and Stanford, including McCarthy (Adelson-Velskii 1970; Feigenbaum 1969). The USSR did not, however, ever achieve influential status in the AI community.

Also reacting to Japan’s Fifth Generation program, the British government established the five-year Alvey Programme in 1983. AI was the “initial focus” of Alvey, though it encompassed other computing research as well (Oakley and Owen 1989, 172). The Programme received £290 million in funding from a combination of public and private sources (Oakley and Owen 1989, 195). It did not posit goals as optimistic and those of Japan or the USA, but whether it constituted a success was a considerable debate. It is held by some in the field to have rejuvenated AI research in the UK, which was largely stagnant since the Lighthill Report of 1973, and also drummed up interest in expert systems (Oakley and Owen 1989, 172). Also in 1983, the Commission of European Communities launched the ten-year European Strategic Program on Research in Information Technology (ESPRIT). This project was designed to develop the industrial potential of European information technology with a strong focus on AI (Steels and



Lepape 1993, 5). It consisted of two phases, the first of which included 32 AI projects funded at cost of 84 million ECU. By 1993, the ongoing second phase involved 27 more projects involving AI with funding around 80 million ECU (Steels and Lepape 1993, 4).

Although each of these strategic computing projects varied considerably in the details of their organization and their particular successes, a few useful generalizations can be made. All were driven by national and supranational bodies which recognized a need for direct intervention in high tech and particularly, AI. All involved forced collaboration across academic, commercial and industrial sectors (and national borders in the case of ESPRIT). All largely failed to achieve their original goals, whether in particular AI achievements or social/economic benefits. However, despite this, all of these projects contributed to strengthening their relative nation's or community's computing and AI infrastructure in terms of products, facilities and skilled developers and researchers.

The embrace of AI, and particularly expert systems, by capital and state during the 1980s generated AI's first critical backlash. A spate of articles wary of the emergent capacities of expert systems emerged. The engineer and militant trade unionist Cooley (1981) notes that with the proliferation of computers, Taylorism can now be applied to intellectual as well as manual work: "fixed capital appropriates not only living labour ... [but also] the scientific and intellectual output of the white-collar workers" (46-47). Positing a need for a forum to discuss the social and political implications of "the deluge of quasi-expert systems [and] artificial-intelligence software tools," Cooley founded the journal *AI and Society*, which still exists today (Cooley 1987, 179). A number of other Marxist voices agreed with Cooley's analysis of AI. Athanasiou (1985) described AI as "cleverly disguised politics". Berman (1992) elaborated this position:

the ideological importance of AI can best be understood as analogous to the role played by scientific management in the second industrial revolution ... AI seeks to achieve the same control over mental processes that scientific management sought to achieve over physical labour through a process of rationalization, fragmentation, mechanization and routinization" (104-105).

Critics of this period did not see their worst fears for AI realized, as by the early 1990s the AI industry went into a decline which it would not recover from for over two decades. This was largely based on the limitations inherent to the expert system approach to AI.

### 3.6 The Decline of Expert Systems

Despite the efforts of numerous states, the AI industry underwent a second AI winter in the early 1990s. Commonly cited reasons for this second AI winter are the “brittleness” of expert systems or their performance dropping to zero when posed a problem beyond their defined scope (Feigenbaum, McCorduck and Nii 1988, 253) as well as their “opacity” or the unpredictable ways they perform as more and more interacting rules are added (Crevier 1993, 204). However, there are other factors as well. The first was a shift in the underlying technology. Expert systems generally ran on LISP machines specially built for the purpose, which cost tens of thousands of dollars. However, in 1977, the Commodore PET (\$795) and the Apple II (\$1298) were released and the personal computing became possible. Cost, ease of use and wide applicability diverted interest from LISP machines, and as PCs became more powerful, AI could run on them as well.

It was not just a technological issue, however. Labour was a problem for expert systems. Some workers resented their knowledge being captured by the knowledge engineering process. Nor was management always enthusiastic towards the large expenditures necessary for hardware and wages of knowledge engineers. One of the creators of XCON noted in 1984 that a “lot of missionary work is still needed” for expert systems to become widely accepted (Kraft 1984, 48). Perhaps more significantly, there turned out to be more labour involved in producing and maintaining expert systems than expected. As expert systems passed from the research lab to commercial application they were applied to increasingly complex domains and so required more and more rules to capture all possibilities. Imagined future expert systems were calculated to require impossibly herculean efforts on the part of knowledge engineers, such that by the mid-1980s, “knowledge engineers were starting to realize that building truly extensive AI systems would require automating the knowledge-acquisition process itself” (Crevier 1993, 205).

Even DARPA recognized that “production of the systems had to be dramatically streamlined” (Roland and Shiman 2002, 194). In addition, the problem of maintenance became apparent. Expert systems needed to be updated every time their problem domains changed. With the constant revolutionizing of business processes required by capitalist competition, the knowledge engineering task for commercial expert systems was never-ending. The solution to both problems was to produce recursive expert systems which could perform their own knowledge engineering, or “to codify for computer use the very knowledge that is helpful in building expert systems” (Feigenbaum, McCorduck and Nii 1988, 250). This proved intractable. Writing in the mid-1980s, proponents of expert systems admitted that “[l]earning is the ‘magic bullet’ that is needed to help with the building of the large knowledge bases” (Feigenbaum, McCorduck and Nii 1988, 255). However, expert systems could not be taught to learn, and by the late 1990s, the industry was in decline. The term “expert system” faded away as the technology was invisibly embedded within more mundane information technologies (Angeli 2010, 52). The search for recursive AI would shift elsewhere.

During this second AI winter, new technologies, such as the internet, took central stage in public interest while AI languished in academic shadows. The late 1990s saw the rise of what Schiller (1999) calls “digital capitalism,” defined by the global market’s reconstitution around the infrastructure of the internet. This period saw a mania of speculative investment in early internet companies. According to one analysis, the release of the Mosaic web browser enabled easy and increased internet access. With it, internet companies proliferated, with investors happy to throw money at anything purporting to do business on the net (Kline 2003). Around 50,000 companies attempted to turn profits on the internet and gathered around \$256 billion in investments (Goldfarb, Kirsch and Pfarrer 2005, 2). Several of these internet startups became today’s AI producers: Amazon, an internet retailer, was founded in 1994 and both Google, initially focused on internet search, and Tencent, initially focused on social media, were both founded in 1998.

As the internet and other information technologies diffused throughout industry, making communication and coordination at a global scale easier, industrial production did not

disappear, but changed. Theorists inspired by the production model developed by Japan's Toyota promoted "lean" production as a successor to Fordist mass production. Lean production aims to have the quality and flexibility of craft production as well as the quantity of mass production, or in other words it "employ[s] teams of multiskilled workers at all levels of the organization and use[s] highly flexible, increasingly automated machines to produce volumes of products in enormous variety" (Womack et al. 1990, 13). While lean production has a spotted track record with less than 10% of companies adopting it doing so successfully (Bhasin and Burcher 2006), it has been widely adopted in software, and AI, production.

The widespread technological optimism of the 1990s took a dramatic blow when the "dot com" speculative investment bubble collapsed in 2000, drawing the global economy into a recession by the following year. Amazon, pre-AI orientation, was devastated with its stock falling from \$107 to just \$7 (Edwards 2016). Many companies, including some large ones such as Pets.com, were killed off in the crash. Others, such as Amazon recovered, in part due to its establishment of the cloud platform Amazon Web Services (AWS) in 2002. Companies such as Google thrived in the post-crash economy and began diversifying their businesses. For instance, Google acquired video-streaming site YouTube in 2006 and online ad company DoubleClick in 2007. AI was no longer a buzzword in the tech industry, but AI research did not cease. It continued, behind the scenes, in academia. GOFAI and expert systems had failed to deliver, but competing AI paradigms were on the rise.<sup>32</sup>

---

<sup>32</sup> One of these competing paradigms was the "situated, embodied, dynamical" (SED) framework (Beer 2014). SED plays a peripheral role in the history of AI and even less in the formation of today's AI Industry. SED refers to a variety of approaches to AI, developed since the 1980s, that emphasize the irreducible importance of the body and its perceptual apparatuses to cognition (Beer 2014, 128). Advocates argue that any attempt to engineer intelligence which only emulates disembodied high-level cognition is doomed to fail when presented with the complexities of concrete reality. The SED approach has thus been concerned with robotics as well as AI. Against the symbolic representations of GOFAI, SED pioneer Rodney Brooks (1991) asserts that it is "better to use the world as its own model" (140). Aiming to circumvent the frame problem, SED systems "do not contain a complicated symbolic model of their environment. Information is left 'out in the world' until such time as the system needs it" (Copeland 2000). Sensors may therefore replace symbolic structures, to some degree. However, as Beer (2014) notes, "abstract reasoning is not rejected by situated approaches, but rather relegated to a supporting role as an evolutionarily recent elaboration of a more basic capacity for getting around in the world" (131). The SED

### 3.7 The Rise of Machine Learning

*“The Industrial Revolution automated manual work and the Information Revolution automated did the same for mental work, but machine learning automates automation itself” (Domingos 2015, 9-10).*

There is an alternate paradigm for AI, which has also existed since the Dartmouth workshop. This is the machine learning (ML) approach. Although machine learning was around since 1956, successful attempts to commercialize it were not made until the 1980s and it did not become commercially widespread until after 2010 when it became the main force of the AI industry’s second era.<sup>33</sup> With machine learning and the new level of recursion it enables, capital would find a new technological base around which to reconfigure its processes.

In 1969, Minsky and Papert (1969) undertook a critical analysis of an early machine learning system called the perceptron. The book demonstrated that perceptrons are unable to process the logical function of exclusive disjunction (XOR). This inability to process a basic logical function generated a lot of skepticism for the machine learning paradigm. While more complex machine learning systems (with hidden layers) were later developed that could implement XOR, it is generally held that Perceptrons caused research into connectionist systems to largely die off until the late 1980s. However, throughout the 1970s, Geoffrey Hinton, David Rumelhart and James McClelland continued to research machine learning at the University of California San Diego. They headed a group called the PDP (parallel distributed processing) group (Nilsson 2010, 339). Their work was not at first generally regarded as being of much importance. Nilsson (2010) notes that “before

---

approach has enjoyed limited commercial success, although humanoid robots based on SED approaches are being developed for industrial settings by companies such as ReThink Robotics (Crowe 2018).

<sup>33</sup> Machine learning is not an equivalent term to artificial neural network. The ANN is one possible way of doing ML. As Mackenzie (2017) demonstrates, there are many different ways to do ML. However, since the various approaches to ML are not relevant here, beyond the types of learning, I use ML and ANN interchangeably. This is also motivated by ANNs, in the form of deep learning, being the cutting-edge in ML today.

about 1980 machine learning (represented mainly by neural network methods) was regarded by some as on the fringes of AI” (398). This attitude was put to rest in 1986 with the PDP group’s publication of *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (Hinton, McClelland and Rumelhart 1986). This collection contained many now-fundamental papers, including one which popularized the “backpropagation” algorithm which was first discussed by Paul Werbos in 1974, enabling many new machine learning applications.

Machine learning takes a very different approach to AI than GOFAI, in at least two respects. The first distinguishing characteristic of machine learning is its statistical nature. The predictions or descriptions made by a machine learning system are probability functions, not logical statements. Thus, the well-known dictum of statisticians that correlation does not equal causation applies to machine learning, although at least one researcher has recently argued that there are ways to overcome this (Hao 2019). Because of its statistical nature, machine learning has not primarily been applied to emulate high-level, usually conscious, human cognitive functions such as logical reasoning. Instead, machine learning has usually been applied to achieve the machinic implementation of lower-level, pattern-recognition functions that may lie partially or completely outside of conscious awareness. This include machine vision, including the recognition of shapes, objects and faces. The statistical nature of machine learning also makes it very useful for finding patterns in vast quantities of data which humans would be unlikely to ever discover.

Second, machine learning systems are distinguished from GOFAI by their capacity to learn, or to “extract patterns from data” (Kaplan 2016, 27). As Alpaydin puts it, to “solve a problem on a computer, we need an algorithm ... For some tasks, however, we do not have an algorithm” (2014, 1). If we cannot create a suitable algorithm, we can use machine learning “to extract automatically the algorithm for [a] task” from available data (Alpaydin 2014, 2). From this recursive automaticity machine learning derives its novel capacities for automation.

Machine learning systems learn through training. While machine learning is an approach that can operate on a variety of architectures, today the cutting edge of machine learning mostly uses artificial neural networks (ANNs)—computer programs that are inspired by (although actually quite different from) the human brain. ANNs roughly mimic the electrical operations of the brain’s neuronal connections rather than emulate high-level logic; they are “based on the assumption that cognition emerges through the interactions of a large number of simple processing elements or units” called neurons (Sun 2014, 109).<sup>34</sup> The neurons are organized into series of layers which are also connected to each other. The lowest level receives input data. Middle levels—also referred to as hidden layers—process data which is sent up from the layers below them. An output layer at the top of the network outputs a solution to a problem. In general, the more layers an ANN has, the more complex patterns it can find and the more complex problems it can solve.

In an ANN, the artificial synapses which connect the layers of artificial neurons are “weighted” with numeric values representing the strength of the connection. Initially these are set randomly. The network “learns” through adjusting the weights of these connections with a learning algorithm. It adjusts them in accord with an inputted dataset, which might be images of faces or audio clips of people saying hello. The network is exposed to many instances of the chosen object(s) and eventually learns to recognize a face or the word hello. This would not, however, be a very powerful attribute of machine learning if these systems could not apply what they have learned to new situations or data not included in the training dataset. In many cases, what matters most is how good the system is at doing that - its “generalization ability” (Alpaydin 2016, 40). Generalization ability depends not only on the training dataset, but also on the model which “defines the template of relationship between the inputs and the output” or how the system uses the training data (Alpaydin 2016, 36). A model has adjustable parameters which are modified in the course of learning (Alpaydin 2016, 37). Alpaydin (2014) sums up the machine learning approach like this: “[w]e have a model defined up to some parameters, and learning is the execution of a computer program to optimize the parameters of the model

---

<sup>34</sup> These networks are not actually comprised of physical mechanical neurons. They are software.

using the training data or past experience. The model may be predictive to make predictions in the future, or descriptive to gain knowledge from data, or both” (3).

The shift described by Alpaydin from knowing a solution to a problem in advance and writing an algorithm to solve it, as with GOFAI, to using a machine learning approach which generates its own solution from data, indicates a qualitative shift in the nature of AI. As Mittelstadt et al. (2016) cautiously put it, “learning capacities grant algorithms some degree of autonomy” (3). More poignantly, Domingos suggests we can “think of machine learning as the inverse of programming” because with machine learning we have “algorithms that make other algorithms ... computers write their own programs so we don’t have to” (2015, 6-7). Machine learning thus represents a new level of technological recursion or “the application of information to itself” (Jordan 2015, 31). While higher-level programming languages abstract from machine code and are thus inherently “automatic,” as Chun (2005) points out, machine learning, at least ideally, abstracts from knowing the solution to the problem one is trying to solve (30). As I will show in Chapter 5, emerging machine learning techniques continue to tend towards an increasing automaticity.

Machine learning first became commercially successful in the late 1980s when it was applied to fields such as speech recognition and optical character recognition for cheque cashing (Crevier 1993, 215-216; Lisboa and Vellido 2000, vii). In 1989, three hundred companies, “most of them startups founded by researchers, competed for [the ANN] market” (Crevier 1993, 216). However, ten years later researchers noted that “[a]pplications of neural networks in ‘real world’ scale are scant. Even though quite a few companies have nurtured neural network prospective studies, not so many make use of neural networks in every-day-business life” (Vellido, Lisboa and Vaughan 1999, 57). That would not happen for another decade.

The 1990s saw the continued sophistication of AI research, but machine learning was not widely commercialized. Experimental autonomous vehicles appeared and enjoyed moderate successes in controlled environments and the Chess-playing system Deep Blue won its first game against chess master Gary Kasparov in 1993. In 1994, the global



market for AI was placed at “about \$900 million, with North America accounting for two-thirds of that total” (Roland and Shiman 2002, 214). Another analysis valued the AI industry as a whole \$363,380,281.69 in 1995 (Liebowitz 1997, 118). In 1998, perhaps the first AI commodity for the domestic sphere appeared in the form of the Furby toy, developed by Tiger Electronics (now Hasbro). The plush owl-hamster hybrid employed simple AI techniques to mimic human speech and associate sounds and enjoyed massive commercial success.

### 3.8 Canada, Deep Learning and the contemporary AI Industry

At this point, the story of the AI industry switches focus to Canada. Canada’s first “high-speed, large-scale” computing facility was built in 1952 at the University of Toronto (Vardalas 2001, 45). As in the USA, the military was involved from the first years of computing research, with the Navy being the most active (Vardalas 2001, 14-43). Canada enjoyed an important, if peripheral role in the earliest decades of AI.<sup>35</sup> Canada’s AI researchers tended to leave the country for better-funded labs, usually in the USA (SCC 1983, 63). But in 1982, The Canadian Institute for Advanced Research (CIFAR) was established in Toronto and in 1984 CIFAR established AI and robotics as its first research area. By 1987, Geoffrey Hinton of the PDP Group was drawn from San Diego to CIFAR to continue his research into machine learning.<sup>36</sup> Hinton and a small group of students continued working on machine learning into the 2000s while most other researchers lost

---

<sup>35</sup> In 1973, the Canadian Society for Computational Studies of Intelligence (CSCSI) (later CAIAC) had its first meeting at The University of Western Ontario in London, Ontario. As one government-funded report from 1983 put it, “Canada cannot boast any major AI laboratories that are privately funded. Most of the expertise is concentrated in university departments in small groups that tend to be underbudgeted and overwhelmed with the need to train graduate students” (SCC 1983, 62).

<sup>36</sup> 1987 also saw the founding of Precarn Inc., an industrial research consortium composed of 34 Canadian corporations from various sectors. Precarn was intended to “conduct and promote long-term, precompetitive research and experimental development in advanced robotics and AI development” (National Research Council Canada 1991, 4). Precarn was influential in coordinating research between academia and industry and was involved in developing a number of commercial applications of AI. Money came in for AI from various sources, including a number of other collaborative organizations and government funding agencies, such as NSERC, which invested over \$3.6 million in AI research in 1989-90 (National Research Council Canada 1991, 5). In 1991, “[o]ver 1000” people in Canada were reported working full-time in the field of AI (National Research Council Canada 1991, 14).

interest: “[b]y the early 2000s, the number of researchers who specialized in neural networks dwindled to fewer than half a dozen” (Allen 2015).

More companies which would later become AI giants appeared at this time. In 2000, Baidu was founded, with a focus on internet search, while Facebook, the social medium to end all others, appeared in 2004. This was the era of “Web 2.0” – applications and websites defined by the proliferation of user-generated content and increased ease of use for casual users (DiNucci 1999, 32). In 2006, Hinton and company demonstrated what came to be known as “deep” machine learning because of the numerous layers which compose the networks, with some networks possessing as many as 1000 (LeCun, Bengio and Hinton 2015, 436-444; He, Zhang, Ren & Sun 2016). The CIFAR researchers described deep learning as using “multiple levels of representation, obtained by composing simple but non-linear modules that each transform the representation at one level (starting with the raw input) into a representation at a higher, slightly more abstract level” (LeCun, Bengio and Hinton 2015, 436). In the years that followed, they would refine the deep learning approach and the approach would garner more and more attention.

At the same time, the financial crisis of 2007-2008, spurred by subprime mortgage lending and consequent housing bubble in the USA, amplified by diverse high-risk lending practices by banks, culminated in the near collapse of the global economy. Enormous government bailouts prevented a total collapse of the banks, although the Great Recession still descended on the world in 2009. While it might seem strange to mention this crisis, based as it was on finance, in the context of the AI industry, it was not unconnected to the dot com crash of 2000 with which several of today’s AI giants were implicated. As Perez (2009) argues, the two crises should rather be seen as a “double bubble” or two episodes of the same story: “[t]he first was based on technological innovation, the second on financial innovation, facilitated, accelerated and made global by information technology and the internet” (802). Amongst the widespread financial fallout after the crisis, hundreds of thousands of jobs were lost, income inequality between the richest and poorest increased substantially and wealth was increasingly concentrated in the richest families (Federal Reserve 2014). It also saw the continuation,

noted by both the OECD (2012) and IMF (2017), of the fall in labour's share of national incomes which began in the 1980s. As of 2019, the economy has yet to fully recover from this double crisis. IMF economists note that “[a]mong the economies that experienced a banking crisis in 2007–08, about 85 percent are still operating at output levels below precrisis trends” (Chen, Mrkaic and Nabar 2018). One analyst adds that while “stocks have beaten all performance records on Wall Street, there has been no corresponding economic recovery” (Bruno 2018).

In 2012, while the world economy was still reeling from the crash, the power of deep learning became widely recognized when Hinton and others demonstrated a record-breaking application of the technique to speech recognition (Hinton et al. 2012). In the next few years, the market for machine learning opened and the AI industry began developing its contemporary form. Three technological factors are generally held to have contributed to deep learning's success (Kelly 2014; T.S. 2016). The first is advancement in learning algorithms, particularly the rehabilitation of the backpropagation algorithm. Second is the availability of big data generated by online activity and the proliferation of mobile devices, since machine learning systems need data on which to be trained. Third is the discovery that graphics processing units (GPUs), formerly devoted to computer gaming, work very well for the massive parallel processing necessary for DL.

Enthusiastic about these technological advances, capital was hungry for AI-producing labour. Most of the researchers at CIFAR spun off their research into startups, many of which were bought up by Chinese and Silicon Valley tech giants. Hinton and two students, Alex Krizhevsky and Ilya Sutskever, formed DNNresearch Inc. in 2012 to market applications of deep learning for image recognition and language processing. DNNresearch quickly won a \$600,000 award from Google before being acquired by the tech giant in 2013. The same year, Hinton was hired as a distinguished researcher by the commercial research team Google Brain.<sup>37</sup> Krizhevsky and Sutskever were also hired by Google. Also in 2013, another of Hinton's student's, Yann Lecun, was hired to lead

---

<sup>37</sup> In 2017, Hinton returned to Toronto to act as Chief Scientific Advisor at the newly founded Vector Institute, a nonprofit aiming to advance Canadian AI capabilities.

Facebook's AI Research Group and Yoshua Bengio, who also worked with Hinton at CIFAR, was hired by IBM in 2015 to help commercialize the Watson supercomputer (Knight 2015). Finally, Ruslan Salahutdinov, who studied under Hinton in Toronto, now teaches at CMU and was hired by Apple to direct its AI research program in 2016.

By 2016, *The Economist* (2016a) could state that “artificial intelligence is finally starting to deliver on its promises”. By the same year, nearly all of the Silicon Valley giants had opened in-house AI research labs and tech CEOs such as Google's Sundar Pichai were proclaiming a shift from “mobile-first” to “AI-first” computing (D'Onfro 2016). Even relatively ancient corporations such as General Electric and IBM became interested in both producing AI and incorporating AI into their production processes (Woyke 2017; Boyle 2017). Around these giant companies both old and new there exists a horde of startup companies, winking in and out of existence, and sometimes being bought up by the large firms.

At the time of writing, in 2019, the AI Industry is flourishing. The products of the AI Industry are being purchased and applied in diverse industries. Businesses “are now adopting different AI technologies to capture benefits such as lower labor costs, increased throughput, enhanced quality and lower downtimes” (Tractica 2019, 11). AI automation is once again a reality and with it comes the “return of the machinery question” (The Economist 2016a). The following chapter examines the political economic dynamics of the AI Industry, but first, I will end this history by examining three factors which influenced the industry's formative days.

The first of these is the platform business model. In the wake of the 2008 financial crisis, many large capitals reconfigured their operations around what Srnicek calls the platform model. All of the AI giants, as we will see in the next chapter, have incorporated the platform model into their operations. Platform capitalism is “centred upon extracting and using a particular kind of raw material: data” (Srnicek 2017a, 39). Platforms are “digital infrastructures that enable two or more groups to interact. They ... position themselves as intermediaries that bring together different users ... More often than not, these platforms also come with a series of tools that enable their users to build their own products,

services, and marketplaces” (Srnicsek 2017a, 43). Through their position as intermediaries, platform capitals are able to extract, control and monetize the data which passes through their platforms. Data thus becomes a resource which platform capitals can appropriate for free. The collection of as much data as possible, argues Srnicsek (2017a), is not an unfortunate side effect of bad management at these companies, but is rather a necessary consequence of the platform business model. While privacy advocates debate how to reform platform capital’s data hunger, others have argued that this drive to data collection cannot be curtailed in a mode of production that amounts most accurately to “surveillance capitalism” (Zuboff 2015).

As the world becomes increasingly digitized and data is continually produced in new fields, platforms engage not only in vertical Fordist integration, but primarily “rhizomatic connections driven by a permanent effort to place themselves in key platform positions” (Srnicsek 2017a, 103). One of these rhizomatic branches is AI – as we have seen, all of the AI giants began business in other fields. Yet machine learning has a particularly special relation to the platform model – it too is necessarily dependent on the collection of data. ML requires extensive data for training. Platforms, with their powerful data collection devices, are thus ideal environments for the production of machine learning. In turn, machine learning offers diverse ways to optimize the functioning of platforms (via microtargeted ads, to take one example). Machine learning and platforms thus form a virtuous data cycle. The platform model is also notable for the far-reaching effects it has had on employment. Platforms have enabled a resurgence of the pre-industrial piecework model of employment. Today this is often called “gig” or “on-demand” work (De Stefano 2015). The next chapter explores how gig work is essential to the AI Industry.

A second factor is the shrinking availability of cheap labour on the global market, particularly in China. Since China’s opening to the world capital market in 1978, capital from the Global North has become accustomed to accessing its abundant supply of cheap manufacturing labour. However, since 2005 the cost of Chinese labour has increased threefold and now exceeds the cost of labour in Brazil, Argentina and Mexico (Gao 2017). This rising wage is at least partially the result of persistent worker struggles in China. Despite dissidence often being met with militaristic force, the China Labour

Bulletin (CLB) notes that worker strikes and protest are on the rise with “more than 1,700 incidents [in 2018], up by 36 percent from 1,250 in 2017” (Symonds 2019).

It is not difficult to see capital’s enthusiasm for AI as, at least in part, a reaction to such wage increases. As the global cost of labour rises, the incentive to AI-powered automation will increase. As the think tank Brookings notes laconically, “[r]ising wages make computers cost-effective for an increasing number of low-skill tasks” (Karsten and West 2015). Already in China, large capitals are feeling the pressure. Foxconn, the electronics manufacturing giant, is seeing substantial profit decreases. Foxconn has already deployed thousands of production robots, but it now reportedly planning a \$4 billion “rollout of cutting-edge robots and also envisions using high-resolution 8K sensors to spot defects that cannot be seen with the human eye” (Ihara 2018). In 2016, President Trump ballyhooed his success in pressuring the heating and air conditioning manufacturer Carrier to cancel its planned move of 1000 jobs from Indiana to Mexico. However, the president of Carrier, Greg Hayes, later admitted that “to continue to be competitive” the company would be investing \$16 million USD in automation, which “ultimately means ... there will be fewer jobs” (Turner 2016).

A third factor is renewed interest in military applications of AI. The post 9/11 USA’s war of terror was already invested heavily in “the cyber” but the revelations of Russian interference in the 2016 USA elections brought even more attention to the diverse ways information technologies might be offensively mobilized (Dyer-Witthford and Matviyenko 2019). The past decade has also seen the military-funded development of diverse weapons incorporating various types of AI. AI a component in autonomous battle systems, vehicles including drones, battlefield robots and cyborg soldier augmentation technology (Scharre 2018; Surber 2018).

In 2017, President Putin declared, “Whoever leads in AI will rule the world” (RT 2017). Other national governments seem to agree. The US Department of Defence formed the JAIC (Joint Artificial Intelligence Center) in 2018 with a mission to “transform the DoD by accelerating the delivery and adoption of AI” (CIO DoD 2018). The top five US defense contractors, including Lockheed Martin, all have AI projects underway (Roth

2019) and many of the major AI producers, such as Google, have been revealed to be doing work for military agencies (to the chagrin of many of their employees). While Google has pledged not to work on offensive weapons, Oberhaus (2018) argues that the interaction between the military and AI is unlikely to cease because AI producers must continue to sell commodities: “once you’ve dominated civilian markets, the capitalist imperative to grow doesn’t just magically stop”.

### 3.9 Conclusion

This chapter has presented a brief political economic history of the AI Industry. It has shown that AI has been implicated with capital, state and military since its inception. AI has not led a purely ideal or conceptual life; it has continually been mobilized as an automation technology. Today, AI research and development it is almost exclusively controlled by capital, as the next chapter, which examines the political economy of the contemporary AI Industry, shows.

## Chapter 4

### 4 Political Economy of the AI Industry

This chapter explores the political economy of today's AI Industry. It aims to situate contemporary AI as the product of a rapidly growing, highly competitive and increasingly concentrated capitalist industry. First, this chapter discusses the capital side of the AI Industry. I survey attempts to quantify the industry, the various types of companies and institutions that make it up and the characteristics that distinguish it from others. These include the involvement of the state, the high degree of concentration and the particular axes of competition, including the cloud computing and AI hardware markets as well as the vigorous open-source AI movement. Second, I discuss the labour side of the AI Industry. I distinguish the different jobs in the industry, from highly paid data scientists to microtask gig workers who earn mere dollars a day. I also discuss how the AI Industry is sharply stratified by both gender and race, with white males occupying the vast majority of lucrative positions. I also discuss how the industry is characterized by rampant sexism and racism and how this has unique consequences for an industry built around machine learning, which relies on quality training data. Finally, I discuss how AI Industry and other tech industry workers have recently bucked a historical trend and begun organizing around such issues as sexism and the militarization of AI research. With this picture of the AI Industry painted, the stage is set for the analysis the machine learning labour process in the following chapter.

#### 4.1 Charting the AI Industry

While firms in diverse industries are now interested in applying AI in their business processes in a variety of ways, there are relatively few companies which actually produce AI technology as a commodity. It is these producers of commoditized AI which I refer to when I mention the AI Industry.<sup>38</sup> The AI Industry spans the globe although it is

---

<sup>38</sup> The producers of AI are not identical to those investing in AI. It is not very surprising to know that the sector which invests the most in AI is the software and information technology sector, with 32% (Naimat 2016, 7). But some unexpected names are also investing large amounts on AI. Naimat (2016) puts the top



unevenly concentrated. The USA is the powerhouse of the AI Industry, followed by China. Other significant countries include Canada, Israel and the UK (Jang 2017; Rapp and O’Keefe 2018).

Attempts at grasping the AI Industry quantitatively have produced varied results. This is due to the industry’s nascent and evolving nature and a lack of consensus on how to define and segment it (Faggella 2019; Zilis and Cham 2016). A quick survey of a few available figures confirms this. One 2016 report states that the “global market for smart machines reached \$6.6 billion in 2015” (McWilliams 2016). Another report published the same year reckons that “the revenue generated from the direct and indirect application of AI software is estimated to grow from \$643.7 million in 2016 to \$36.8 billion by 2025” (Tractica 2016). More recently, the same source puts the global AI market at \$3.2 billion USD in 2016 and expects it will grow to \$89.85 billion USD by 2025 (Tractica 2019). International Data Corporation (2018) places global spending on AI at \$19.1 billion in 2018 and predicts it will reach \$52.2 billion by 2021. The one thing these numbers confirm is that capital is excited about the prospects of the AI Industry. It is certain that the industry is growing rapidly.

In contrast to the first manifestation of the AI Industry, which revolved solely around GOFAI expert systems, today’s burgeoning AI Industry is based on machine learning, which has myriad potential applications.<sup>39</sup> However, it is important to note that the majority of existing AI companies are still in early stages of development. 62% of the 1548 AI companies surveyed by Naimat (2016) were still in the “lab project” stage while

---

18 investors in AI in order as: Google, Facebook, Rocket Fuel, IBM, Amazon, Yahoo, Intel, Microsoft, Deloitte, MITRE, Baidu, LinkedIn, Apple, Cylance, Lockheed Martin, NASA, Sentient Corporation, Electronic Arts (8). Deloitte is a multinational accounting and consulting firm, while MITRE is a non-profit research and development organization in the USA and Electronic Arts is a video game company.

<sup>39</sup> According to Dong (2017), while AI research labs at the tech giants focus on cutting-edge deep learning “most applications of machine learning at these same companies do not rely on neural networks and instead use traditional machine learning models. The most common models include linear/logistic regression, random forests and boosted decision trees. These are the models behind ... friend suggestions, ad targeting, user interest prediction, supply/demand simulation and search result ranking”.

a mere 5.6% were ready for “strategic direction for business” (6).<sup>40</sup> Despite this, it is worth surveying the variety of AI commodities under development. Rather than discuss particular AI products, I will survey some of the attempts that have been made to segment the industry by type of product.

One approach is to divide by the type of AI subfield employed. Chen (2016) suggests the following breakdown: logical reasoning, knowledge representation, planning and navigation, natural language processing, and perception. According to a similarly framed study, just over a third (505) of a total 1500 AI startups are working on natural language processing, with image processing coming in second at 197. With 169 is cognitive computing (IBM’s personal term for AI), with 98 is autonomous vehicles, and with 58, chatbots (Naimat 2016, 14). Research from Tractica takes a different approach, dividing the industry into three broad subfield groups: big data, vision and language, with big data representing the majority of uses cases in applications such as analytics, data processing and algorithmic trading, though vision and language are gaining ground (Groopman 2017, 4). Another market research firm describes the industry in terms of six applications: virtual reality assistants, intelligent agents, expert systems, embedded software, autonomous robots/vehicles and purpose-built smart machines (BCC Research 2016). Another (ambitious) attempt has been made, by Zilis and Cham (2016), to segment the AI Industry by field of application. Only a few of these include education, logistics, customer support, security, agriculture, medicine, as well as a variety of enterprise and technology applications. According to one study of the 50 highest-funded AI startups, the most popular fields were advertising and customer resource management, followed by core AI technologies and business intelligence and analytics (Pham 2017). Another meta-analysis claims that healthcare, finance and marketing appear to be the most popular applications (Faggella 2019).

---

<sup>40</sup> It is also important to note that most successful applications are based on relatively simple ML techniques rather than the cutting edge. According to a CEO I spoke to, “there’s an increasing mismatch between the requirements of industry and the progress in academia. For example, a lot of these fancy neural network approaches like adversarial networks and other generative models, nobody wants them. Businesses ... need like text processing, speech recognition, they need like the bottom floor. Not the 400th floor” (P 13).

More simply, Brynjolfsson and McAfee (2017) present a useful heuristic division of the AI Industry into two branches which have enjoyed the most successful commercial application: perception and cognition. Perception is the machinic implementation of the functions of the human sensory apparatus. It includes techniques such as machine vision, natural language processing and sentiment analysis. AI perception has been commercially applied in fields such as image recognition, speech recognition and medical diagnosis. As of 2018, many such systems now match or exceed human performance in their narrow domains. Perception is an essential component of autonomous vehicle research. Cognition, in the AI Industry context, has two prevalent forms: classification and prediction. Classification is perhaps the most popular commercial application of machine learning because there are many situations in which a business wants to know whether  $X$  belongs to category A, B or C (Brynjolfsson and McAfee 2017). Prediction can be thought of as extrapolating from classifications. The credit agency Equifax claims to have applied machine learning to improve the predictive ability of their credit scoring models by 15% (Press 2017). Some analysts expect that prediction will ultimately become the characteristic commercial application of machine learning (Agrawal, Gans & Goldfarb 2016).

High-level portraits of the AI industry as a whole, while giving a useful orientation to the field, necessarily lack detail as to the industrial organization of its various branches. Therefore, the next section discusses the five major types of organizations involved. These include three types of companies: first, a handful of giant tech capitals hailing from the USA and China. These include Facebook, Google, Amazon, IBM and Microsoft and China's Baidu and Tencent. Second, are large, older corporate "dinosaurs" such as General Electric and Siemens, which are attempting to modernize by both incorporating AI into their production processes and moving into AI production. Third, is a globally distributed horde of startup companies which are often bought up by the larger firms. Besides these three types of companies, there exist a number of think tanks and research laboratories. Some of these operate as non-profits while others are funded by a combination of academia, industry and government. Indeed, state involvement in AI production is significant enough that it constitutes the fifth type of involved organization.

In the next section, I analyse examples of each type of organization and show how each exhibits the typical characteristics of such an organization.

## 4.2 AI Tech Giants

As early as 2016, *The Economist* (2016b) noted that in “Silicon Valley a handful of giants are enjoying market shares and profit margins not seen since the robber barons in the late 19th century”. Today these powerful tech capitals have nearly all delved into the production of AI, which as we have seen, has a recursive relationship with the datacentric nature of platforms. According to Krauth (2018) the top investors (in the USA) in AI research are:

1. Google - \$3.9 billion
2. Amazon - \$871 million
3. Apple - \$786 million
4. Intel - \$776 million
5. Microsoft - \$690 million
6. Uber - \$680 million
7. Twitter - \$629 million
8. AOL - \$191.7 million
9. Facebook - \$60 million
10. Salesforce - \$32.8 million

Unsurprisingly, the USA AI tech giants are well represented in this list. Also unsurprisingly, all of these companies incorporate the platform model discussed in the previous chapter. We have already noted the centrality of data to these companies, and the rhizomatic expansion into new markets and data sources that this entails, but another characteristic is worth noting. Srnicek (2017a) holds that, driven by data hunger, platforms are compelled not only to integrate vertically or horizontally like older businesses, but instead attempt to occupy key positions wherever data may be extracted (106-107).<sup>41</sup> This leads to what Srnicek (2017a) terms the “convergence thesis” or the theory that there is a “tendency for different platform companies to become increasingly similar as they encroach upon the same [key] market and data areas” (107). Thus, to take only a few examples, Google the search company has attempted to create a social media

---

<sup>41</sup> For a detailed study of how contemporary cybernetic capitalism is built upon a complex “stack” of interacting layers of technological infrastructure see Bratton (2016).

platform (Google+) while Facebook, the social media platform, has attempted to move into online retail, which Amazon dominates. More recently, Facebook has announced that it will release its own cryptocurrency called Libra, which it hopes will allow it control over payment data (Costine 2019). Whether Srnicek's convergence thesis proves completely true or not, AI is one case where it does seem to be verified. The tech giants are all tending towards a platform or cloud AI model.

I will explore typical features of AI tech giants by looking at the example of Google, even if Google is something of a special case because it dwarfs the other tech giants in scale and diversity. Founded in 1998, Google weathered the dot com crash of 2000-2002 and rapidly ascended the corporate technology hierarchy thereafter. Based on the Google search algorithm, the company now owns the most popular search engine in the world, the biggest streaming video platform (YouTube), the most popular web browser (Chrome) and email service (Gmail) as well as the most widely used mobile operating system (Android) (Dickey 2014). Google first publicly delved into AI proper with the founding of Google Brain in 2011. Shortly after, Google Brain researchers demonstrated a deep learning system that learned, via unsupervised learning, to recognize cats in YouTube videos (Markoff 2012). This feat was followed by a spate of cutting-edge AI research, much of which Google researchers share through publishing in academic journals as well as the Google AI Blog.

Like most of the AI tech giants, Google produces AI for both the consumer commodity (nonproductive use) and fixed capital (used in production processes) markets. Google's most visible AI consumer commodity is the Google Home "smart speaker" first produced in 2016. This is a physical platform for Google Assistant which employs various types of AI, including speech recognition, to perform entertainment, organization and home automation tasks via spoken word commands. Ongoing incidents have revealed that Home and other smart speakers also function as data gathering devices (Day 2019).

Google's most significant AI products are, however, sold for use as fixed capital by other capitals. This occurs primarily through Google Cloud Platform, which at \$1 billion per quarter in 2017 is the fastest growing business in the company (Trefis Team 2019). While

the cloud was originally limited to providing storage functions, it is now used to provide computing and AI functions. Computing power necessary for cutting-edge research continues to increase. Amodei and Hernandez (2018) calculate that since 2012, the amount of computing power “used in the largest AI training runs has been increasing exponentially with a 3.5 month doubling time”.<sup>42</sup> Since few companies can afford to own expensive computing hardware they obtain access to it via the cloud.

Google’s Cloud AI offering include a platform for AI developers to train models as well as a variety of tools, applications, “plug and play AI components” and “building blocks” which “make it easy for developers to add sight, language, conversation, and structured data to their applications” (Google n.d.a). These offerings help data scientists and other skilled AI workers integrate existing machine learning models into their businesses. Google also offers “prepackaged” AI solutions to common business problems targeted at users lacking expertise in AI and machine learning. The “Document Understanding AI” automates the reading and processing of documents while the “Contact Center AI” automates customer service via speech or text. These solutions are purported to “deliver personalized customer experiences, increase sales while lowering costs, and get the insights you need to make better business decisions” (Google n.d.b). Finally, as I will discuss in greater detail in Chapter 5, Google has since 2018, been offering an “automated way for businesses to build new algorithms” (Metz 2018b). This automated machine learning aims to make machine learning available to businesses lacking technical expertise. Several of the other AI tech giants are converging around similar cloud AI platforms. Amazon Web Services is the dominant platform with 35% market share (Miller 2017). Other cloud platforms include Microsoft Azure, IBM Cloud and Alibaba Cloud. Facebook has yet to enter the cloud AI market.

---

<sup>42</sup> The energy consumption of cloud facilities is staggering. In 2014, data centers in the USA used 70 billion kWh, about 1.8% of total US energy consumption (Shehabi et al 2016, ES-1). While the largest of these are operating increasingly efficiently in terms of energy consumption in recent years, there is no upper limit on how large or how numerous such centers will become (Shehabi et al 2016, ES-1). It has been predicted that Iceland will soon expend more energy powering cryptocurrency mining data centers than it does on residential consumption (Farivar 2018).

Beyond its AI commodity and fixed capital production, Google is engaged in “moonshot” research. Some of these, like Google Glass, have failed spectacularly. Others, especially in AI, have paid off remarkably. In 2014, Google acquired the UK startup DeepMind, a “startup with no revenue or marketable product but a team of ‘deep learning’” researchers” for £400 million (The Economist 2017). Since then DeepMind has gone on to produce groundbreaking machine learning research such as AlphaZero. DeepMind’s (n.d.a) stated goal is to “Solve Intelligence. Use it to make the world a better place”. They describe their research as striving towards the creation of AGI, or in other words, “developing programs that can learn to solve any complex problem without needing to be taught how” (DeepMind n.d.b). While this remains a distant goal, Google is already actively redirecting DeepMind’s work towards commercializable healthcare applications (DeepMind n.d.c).

Yet, while AI is now a central interest for Google and other tech giants like Amazon and Facebook, none of these companies is solely invested in AI. For instance, Google generates around 85% of its overall revenue from advertising (Schomer 2019). However, Google is continually expanding its range of AI products. It also recursively integrates its AI products into its own operations. AI is central to its advertising efforts and a DeepMind AI was applied to reduce the energy usage of a Google data center by over 40% (Evans and Gao 2016). AI is also integral to the ostensibly free product-services which such companies provide to users in (tacit) exchange for access to data. Google Maps and Translate both employ AI to perform their respective functions of navigation and language translation. The data gathered from users of these services is then fed back into them to improve their operation, but is also added to Google’s larger data vaults.

In sum, Google is typical of the AI giants insofar as it increasingly integrates AI functions into its diverse business involvements, most or all of which give Google access to vast amounts of data, which it uses to power its AI efforts. Like most AI giants, Google produces for both commodity and fixed capital markets, and delivers much of its fixed capital AI via a cloud platform (Facebook has yet to enter the cloud computing or cloud AI market). Like all other AI giants, Google is also engaged in the frequent acquisition of AI startups and continually strives to break into new markets.

While an ideal representative of its type, Google is distinguished from the other giants by its sheer size, diversified nature and massive investment in research. It is nearly impossible to use the internet without coming into contact with a Google product or service. This pervasiveness makes it unique, at least beyond China, where Baidu controls most of internet search and Baidu and Tencent dominate video streaming (Sun 2018).

### 4.3 AI Dinosaurs

General Electric (GE), founded in 1892 in Boston, Massachusetts, is a different breed of AI company. Originally a producer of a variety of electrical products, GE was also a contributor to the early computing industry of the 1960s. Today GE is a huge conglomerate and is invested in diverse fields from healthcare to aviation. Around 2015, GE made a move to modernize its operations by focusing on AI, machine learning and other 4<sup>th</sup> industrial revolution technologies (Woyke 2017). GE (n.d.a) offers a wide range of research, design and fabrication services which are now augmented by an AI-powered platform which offers capacities for simulation, testing and prediction.

However, the main product of AI dinosaurs is fixed capital for other industrial producers. GE (nd.b) sells AI-powered software for industrial automation, such as iFIX, which allows a plant operator to see and control all of the relevant operations of a factory in one app, thereby “[r]educ[ing] costs and risk”. GE is also selling the Predix Industrial IoT Platform which it is marketed as an operating system for factories (Passieri 2015). Predix integrates data from diverse aspects of production and circulation processes and allows for a predictive rather than reactive approach to managing these processes. It is described as providing “the software architecture and services required to make any machine an intelligent asset” (Predix n.d.). The even older German conglomerate Siemens (founded 1847) offers a similar product called MindSphere, which it describes as a “cloud-based, open IoT operating system ... that connects your products, plants, systems, and machines” (Siemens n.d.).

AI dinosaurs are, like tech giants, eager to buy up startups. Siemens spent \$10 billion USD on acquiring software startups from 2007 to 2017 (Walker 2019). While AI dinosaurs do buy technology from the AI giants if necessary (Siemens uses IBM’s



Watson Analytics) they are generally attempting to develop their own AI in-house. According to one analysis, “most of the major companies making the machine learning tools for manufacturing are also using the same tools in their own manufacturing. This makes them the developer, the test case and the first customers for many of these advances” (Walker 2019).

While dinosaurs like GE and Siemens are focusing on machine learning-driven analytics, others like Fanuc and KUKA, are using AI to power smart robots (Walker 2019). In different ways, AI dinosaurs are providing new ways for industrial capitals to “bring down labor costs, reduce product defects, shorten unplanned downtimes, improve transition times, and increase production speed” (Walker 2019). Whether or not these companies will converge with the AI tech giants remains to be seen, though it currently seems unlikely GE will delve into social media any time soon.

#### 4.4 AI Startups

AI startups are smaller companies which often rely on venture capital to fund their operations. The startup companies that I interviewed workers from ranged in size from only six employees to more than thirty. AI startups operate a variety of business models. One machine learning scientist told me that startup companies deliver premade models, insights or services. Recently, startups have even begun to produce specialized AI hardware (Metz 2018c).

Premade AI models are made for diverse usages. One of my interviewees worked on a voice biometrics system for jails. Other interviewees described working on a variety of marketing and advertising applications requiring predictive analytics, while another was applying machine vision to healthcare. It is impossible to neatly sum up the types of commodities produced by AI startups. Companies which provide insights or services operate a consulting business model or build customized AI systems for the particular needs of other businesses. Again, the potential applications vary widely. One of my interviewees did a lot of AI consulting work in the financial sector, while another had recently been advising on a natural language processing system for incorporation in smart televisions.

One recent analysis puts the global total number of AI startups at 3465 (Fabian 2018). Fabian (2018) found 1393 startups located in the USA, China following with 383 and Canada in fifth place with 131. Naimat (2016) reckons that “1,500 companies in North America that are doing anything related to AI today” though only 87 have projects mature enough for deployment in business (8-13). Most AI startups depend on venture capital funding to get off the ground. Benaich (2016) holds that in 2015, 267 AI firms raised \$1.9 [billion] in venture capital funding [in North America]. AI Index (2018) counted just over 750 VC funded AI startups in the US in 2018, an increase of 113% from 2015 (31). However, the number of AI startups at any moment is difficult to place with certainty since many simply die off while also there is a “massive amount of venture funding and acquisition”; 115 out of 120 AI startups that exited the market in 2017 were acquired (Patrizio 2018). *The Economist* (2017) noted \$21.3 billion in mergers and acquisitions related to AI in 2017. Google alone has acquired at least 30 AI startups since 2014. AI startups are generally closely linked to the AI tech giants even before being bought up. Many receive funding before being acquired and almost all AI startups purchase cloud computing power from the AI giants and employ their AI tools.

In 2018, Chinese startup SenseTime received a \$600 million USD round of funding, led by Alibaba, which put its total valuation at more than \$4.5 billion USD, making it the most valuable AI startup in the world at the time (Vincent 2018). Founded in 2014, the company became profitable in 2017 and claims to have “more than 400 clients and partners” (Vincent 2018). Its products focus on machine vision and range from smartphone “beautification” filters for social media to autonomous vehicles to automated surveillance video analysis sold to police (Vincent 2018). This is one path AI startups may take. Others, like DeepMind, as discussed above, aim to be acquired before ever producing a working product.

## 4.5 AI Think Tanks

Think tanks are private research institutes that may be non-profit or funded by academia, industry or advocacy groups. While AI has been involved with think tanks since its inception (i.e. RAND), think tanks devoted specifically to AI only began to appear in the mid 2010s (Think Tank Watch 2018). The nature of these organizations vary. Some

were founded by transhumanists concerned with the dangers and prospects of future AGI and/or ASI. These include the Machine Intelligence Research Institute (MIRI) a non-profit funded by philanthropic groups including the Thiel Foundation, and the Future of Humanity Institute (FHI) at Oxford University. These thinktanks are part of the “AI safety” movement which aims to ensure that if human-level (or higher) AI appears, it can be controlled (Future of Life Institute n.d.). They fund outside academic researchers as well as produce their own research.

Other AI think tanks are more concerned with actually-existing AI and the AI Industry. These usually have connections to both industry and academia, such as the AI Now Institute (part of NYU and funded by Google and Microsoft) and the Partnership on AI (partners include MIT Media Lab, Amazon, Facebook and Google). AI Now (n.d.) describes itself as studying the “social implications of artificial intelligence” while the Partnership on AI (n.d.) states that its goal is to “study and formulate best practices on AI technologies, to advance the public’s understanding of AI, and to serve as an open platform for discussion and engagement about AI and its influences on people and society”. These think tanks do indeed produce valuable research on diverse aspects of the social implications of AI, from gender to automation of work; some of which I draw on in this dissertation. But they serve other functions as well.

One important function is serving as a vehicle for the so-called “democratization” of AI programs which the tech giants are advancing. This will be discussed below in more detail, but for now it is enough to note that these socially-concerned think tanks, and other similar initiatives by AI capitals, play an important public relations function. This came to light definitively in the wake of revelations about Google’s secret agreement to provide machine learning image recognition for the Pentagon’s “Project Maven”. In emails leaked to *The Intercept*, Fei Fei Li, Google’s then head scientist at Google Cloud, worried that the exposure of Project Maven, with which she was centrally involved, could

counteract her company's efforts to portray their AI products as socially benevolent.<sup>43</sup> Li writes:

I don't know what would happen if the media starts picking up a theme that Google is secretly building AI weapons or AI technologies to enable weapons for the Defense industry ... Google Cloud has been building our theme on Democratizing AI in 2017, and Diane and I have been talking about Humanistic AI for enterprise. I'd be super careful to protect these very positive images (quoted in Fang 2018).

It would be incorrect to attribute solely public relations function to AI think tanks, but it would be simply false to deny its importance as AI tech giants come under increasing scrutiny for diverse reasons. I return to this in section on labour. However, there are some other AI think tanks with a different function. One particularly interesting case is that of OpenAI.

OpenAI was founded in 2015 by Elon Musk with \$1 billion USD funding from Musk, Peter Thiel and other business magnates. Sutskever, former student of Hinton, acts as director of research. While, like MIRI and FLI, OpenAI conducts research aimed towards producing safe AGI, it is distinctive in that it actively produces cutting-edge AI systems, such as OpenAI Five, which has defeated skilled players at the video game DOTA2 (OpenAI 2018). While initially operating as a non-profit, in 2019 OpenAI announced a shift to a "capped-profit" model which will allow them to raise the capital necessary to compete with the AI tech giants. They describe the shift as motivated by the "need to invest billions of dollars in upcoming years into large-scale cloud compute, attracting and retaining talented people, and building AI supercomputers" (OpenAI 2019). Since then, Microsoft has invested \$1 billion USD in OpenAI and has become its exclusive cloud provider (Brockman 2019). This shift sums up the difficulties of AI research and development beyond the reach of capital – something I will return to in the final

---

<sup>43</sup> Li is also co-founder of the non-profit AI4ALL (n.d.) which purports to be "dedicated to increasing diversity and inclusion in AI education, research, development, and policy".

chapter. Whether non-profit thinktanks like MIRI will be able to continue to exist without massive support from and integration with AI capital remains to be seen.

## 4.6 AI and the Contemporary State

I showed in the previous chapter how national governments have played key roles in the history of the AI Industry. Nearly forty years after Japan's Fifth Generation Computing initiative and the USA's Strategic Computing Program ended, AI is once again a topic of national interest. Whether or not it is true, there is an increasingly widespread belief that, in the words of China's State Council, AI "has become a new engine of economic development" (Bloomberg News 2017). The USA had the first national government to explicitly engage with the AI Industry. In 2016, the outgoing Obama administration directed the Executive Office of the President to produce two reports on AI. The first report states that the USA government spent approximately \$1.1 billion on unclassified R&D for "AI-related" technologies in 2015 (Executive Office of the President 2016, 25). The report asserts that a "strong case can be made in favor of increased Federal funding" to the tune of doubling or even tripling, with a lucrative expected payoff in economic growth (Executive Office of the President 2016, 25). The report also notes a deficit in qualified AI workers and recommends increased funding for education. In 2019, President Trump signed an executive order to launch the American Artificial Intelligence Initiative which outlines five areas of focus: "research and development, availability of data and resources, ethical standards and governance, education, and international collaboration that also protects American interests" (Wei-Haas 2019). The Initiative is driven by a recognition that "Artificial Intelligence will affect the missions of nearly all executive departments and agencies" (Trump 2019).

In Canada, a despairing report by the Information and Communications Technology Council (2015) described Canada as suffering a dire lack of AI graduate students and funding to attract and keep them in the country. In early 2017, Hinton and others published a piece in *The Globe and Mail* expressing their dissatisfaction with the state of AI in Canada. Their machinations proved effective and later that year, Canada became the first country to launch a fully-funded AI strategic plan (\$125 million), recognizing that AI "has the potential to drive strong economic growth, by improving the way we

produce goods, deliver services and tackle challenges like climate change” (Canadian Federal Government 2017, 103). As part of this strategy, AI research institutes were established in both Toronto and Montreal.

Other states have followed suit. President Macron stated in March 2018 that the French government would spend €1.5 billion over five years to advance national AI research. In the UK, The Alan Turing Institute encourages collaboration between academic and industry researchers with the aim of increasing national AI capacities. In addition, a coalition of researchers from across Europe have written a manifesto which calls for the establishment of a European Lab for Learning & Intelligent Systems (ELLIS) by which to develop AI talent in Europe. At least 18 countries have now announced national AI strategies (Dutton, Barron and Boskovic 2018). Perhaps most significant of these is that of China.

In 2017, China’s State Council issued a notice outlining its Next Generation Artificial Intelligence Development Plan which aims to achieve, by 2020, AI capacities on par with the USA. By 2030, China aims to be the number one producer of AI with a “gross output of RMB 1 trillion (U.S. \$150.8 billion) for the core AI industry and RMB 10 trillion (1.5 trillion) for related industries” (Ding 2018, 7). This plan is significant not only for its unprecedented levels of spending, but also because China is currently the only country with the expertise and resources to contest USA AI dominance. More than one analysis suggests that countries that fail to invest in AI now will fall economically behind forever, as recursive AI productivity gains compel exponential increases in the economies of China and the USA (Cummings et al. 2018 vi; K.F. Lee 2018, 185-186). While K.F. Lee (2018) is reluctant to describe this situation in terms of a national “AI race” because of its zero-sum connotations (246). However, AI is in fact a central component of a new global arms race (Simonite 2017c; Allen 2019).

The Chinese People’s Liberation Army intends a transformation towards AI-powered “intelligentized” warfare (Kania 2017) while in the USA, DARPA is working on COMPASS (Collection and Monitoring via Planning for Active Situational Scenarios) which will “leverage advanced AI and other technologies to help commanders make

more effective decisions to thwart an enemy’s complex, multi-layered disruptive activity” (South 2018). While fully autonomous weapons do not yet exist, or at least remain secret, increasing degrees of semi-autonomy have been achieved (Scharre 2018). Autonomous weapons are, however, only one way that AI might be weaponized. Autonomous weapons threaten physical security, but AI might also be used to threaten digital security via cyberattacks, and political security via surveillance, persuasion and deception (Brundage et al. 2018, 6). Further, AI might exacerbate the risk of nuclear war (Geist and Lohn 2018).

Despite this state interest, Cummings (2018) notes that due to the tech industry’s massive salaries, governments are finding it hard to hire AI experts and are forced to contract out for their AI needs. He wonders if “defence companies and governments continue down a path of relative AI illiteracy, could this enable a potential power shift such that critical AI services will be leased via Google, Amazon or Facebook?” (Cummings 2018, 16). While the possibility of national militaries powered by corporate AI lies beyond the scope of this dissertation, based on the situation of the AI industry today, there is no doubt that AI producers aim to keep concentrated control of the technology they produce.

## 4.7 AI Industry Concentration

An arms race seems like an apt descriptor of the dynamics of the AI Industry. Despite a hangover of early internet era language rife with connotations of freedom, decentralization and democratization, today’s AI industry is “dominated by a small number of internationally active companies and is characterized by a strong trend toward market concentration in all major segments” (Dolata 2018, 102). Strong concentration is enabled by “extraordinary financial strength” which allows the tech giants to invest heavily in continual research, development and fixed capital improvement, making it “very hard for newcomers to become serious competitors of the established leaders in any of the already occupied core business fields” (Dolata 2018, 91). Only the rival giants will be able to engage in “fierce oligopolistic competition ... carried out primarily through aggressive innovation and expansion strategies” (Dolata 2018, 98). *The Economist* (2017) predicts that:

Over the next several years, large tech firms are going to go head-to-head in three ways. They will continue to compete for talent to help train their corporate “brains”; they will try to apply machine learning to their existing businesses more effectively than rivals; and they will try to create new profit centres with the help of AI.

My interviewees also noted the fierce competition of their industry. One data scientist told me “we’re not living in real free market condition. There are bigger forces that dominate the market all the time ... the fate of the industry is tied to ... Google, Amazon ... IBM. Giants ... they hire most of the advanced graduates” (P 9). A startup CEO expressed a similar view: “We’re financially competing with Facebook and Google for talent ... you can’t find the right people because Google and Facebook just bought them all” (P 13). Dolata (2018) expects that this concentrated competition of elites will produce a “remarkable volatility of acquired market and power positions, which must be repeatedly defended and renewed in the face of the extremely rapid succession of innovation dynamics” (103). Several other analyses have come to similar conclusions. *The Economist* (2017) deems it

likely that the incumbent tech groups will capture many of AI’s gains, given their wealth of data, computing power, smart algorithms and human talent, not to mention a head start on investing. History points to the likelihood of concentration; both databases and personal computers ushered in ascendancies, if only for a while, of a tiny group of tech firms.

Data is of particular importance. Since machine learning requires massive, quality datasets, those AI companies which have access to large amounts of data from their operations (e.g. social media, online retail, search) have been the ones to thrive. Agrawal, Gans and Goldfarb (2018) predict further concentration for the AI Industry based on this recursive dynamic of machine learning and data:



AI has scale economies ... AI tools are often characterized by some degree of increasing returns: better prediction accuracy leads to more users, more users generate more data, and more data leads to better prediction accuracy. Businesses have greater incentives to build [AI] if they have more control, but along with scale economies, this may lead to monopolization (23).

K.F. Lee (2018) argues similarly that the “positive-feedback loop generated by increasing amounts of data means that AI-driven industries naturally tend toward monopoly, simultaneously driving down prices and eliminating competition among firms” (161). Practically, this means that companies such as Facebook and Google, which have access to vast, continually accruing quantities of search and social media data have a decided advantage in the industry. One analysis of IBM asserts that because the company does not have data resources like those of Facebook and Google it will necessarily find it difficult to form a “virtuous circle” of AI commercialization in which “launching a product gets users, users generate data, and the data improves the product” (Kisner, Wishnow and Ivannikov 2017, 19-20). Or as one startup CEO put it to me: “large companies that have large amounts of data ... have a huge defensive moat ... it becomes really hard for a startup to come along and disrupt the big one if they don’t have a dataset ... how do they get that dataset if they don’t have any algorithms ... or a deployed product” (P 3). In sum, the “key processes and categories with which to adequately describe the essential developmental trends” of the AI Industry are not “decentralization, democratization and cooperation but rather concentration, control and power” (Dolata 2018, 86). Countervailing forces may, however, be on the horizon.

There is a growing consensus that the tech giants have accrued too much power, even if there is no consensus on how to deal with it. At the time of writing, the US Justice Department and Federal Trade Commission are engaged in an antitrust investigation in the tech sector (Kelly 2019). Some commentators argue that all that is needed is to stimulate competition (Doctorow 2019; *The Economist* 2019). Others call for new types of regulation (Chen 2019; Cath et al. 2018; Reed 2018). Still others, including 2020 US Democrat president candidate Elizabeth Warren (2019) argue for breaking up the tech

giants so that they “do not crowd out potential competitors, smother the next generation of great tech companies, and wield so much power that they can undermine our democracy”. Warren (2019) suggests that the platforms established by the tech giants be split from their other business ventures and operated as “platform utilities”. Precedent for all of this is provided by the EU’s 2018 implementation of the General Data Protection Regulation (GDPR), which puts strict rules on how companies must store and use user data. In January 2019, France fined Google €50 million for violating the GDPR. Shortly after, Apple CEO Tim Cook (2019) published an opinion piece in *Time* calling for similar regulations in the USA. How these discussions will play out remains uncertain. It is hard to guess, for instance, exactly how states will balance the emerging desire to reign in the AI tech giants with the competitive dynamics of an international AI arms race driven by those same companies.

## 4.8 Open Source AI, Clouds, AI Chips

Amongst this turbulence, an interesting phenomenon of open source AI tools has swept across the industry. In 2015, Google open sourced its Tensorflow machine learning software library.<sup>44</sup> Google CEO Sundar Pichai (2015) said he hoped it would help “exchange ideas much more quickly, through working code rather than just research papers. And that, in turn, will accelerate research on machine learning, in the end making technology work better for everyone”. Since then, nearly all of the AI giants have released some of their AI tools as open source, including IBM’s SystemML, Facebook’s PyTorch, Amazon’s Neo-AI and Baidu’s Warp-CTC. According to my interviews, use of these open source tools is near universal. One CEO I spoke to said, “[m]ost startups ... use open source technologies. Everything from the programming language that we use, which is primarily Python and Javascript ... to the software packages we use” (P 3).

---

<sup>44</sup> Tensorflow was actually not the first open source ML toolkit; pre-existing Google’s open source offerings were at least two other open source libraries suitable for ML, including Theano (developed by the Montreal Institute for Learning Algorithms) and Keras (privately developed primarily by François Chollet, a Google engineer). The open source programs of Google and the other AI giants should be read as a reaction to these predecessors which were produced outside of their control.

Another CEO agreed: “all the tools that we actually use for building the product ... besides [operating systems], are open source. The couple exceptions, like NVIDIA, have some software called cuDNN, that’s proprietary. But every machine learning researcher uses it” (P 2). Several companies also used MathWorks’ MATLAB, a proprietary software suite, but a vast majority of the tools for producing machine learning are freely available.

Why would an industry characterized by aggressive oligopolistic competition be so eager to open source proprietary software? According to the tech giants, it is out of a benevolent desire to bring AI to the masses. Concerning its “Democratizing AI” program, Microsoft (2016) asserts that “We’re going to infuse every application that we interact with, on any device, at any point in time, with intelligence,” while Intel extols the wonders of “ubiquitous artificial intelligence” (Intel Brandvoice 2018). Microsoft (2016) says it will make “these same intelligent capabilities that are infused in our own apps ... available to every application developer in the world”. The goal is “helping everyone achieve more — humans and machines working together to make the world a better place” (Microsoft 2016). Commentators have been more prosaic. Aside from the free contributions of a community of contributing developers, Gershgorn (2015) suggests that such companies “might indirectly benefit from open-sourcing their tools” by both providing skilled employees for future hiring in a tight market and becoming the ground on which future applications are built, thus ensuring the ongoing relevance of their products. This is the same strategy which Google deployed by open sourcing its Android mobile operating system, which is now the most popular open source mobile OS in the world (Gershgorn 2015; Amadeo 2018).

This is, however, only half the story of the corporate assimilation of open source AI, which should be considered alongside two other avenues which many of the AI giants have begun exploring in the last five years: specialized AI hardware and the cloud. In the words of one data scientist I spoke to: “Most of the computation in data science is done on the cloud. We have the local servers also running, but most of the heavyweight data solutions [are] on the cloud” (P 9). Google and Amazon’s cloud services were the most popular, but Microsoft Azure was also mentioned a few times, while a smaller firm, Digital Ocean, received one mention. AI startups do not only use the GPU computing

power of the cloud. Many also employ the premade AI solutions offered by the giants. One CEO told me: “I use Google Cloud a lot ... I use them for speech-to-text ... I use them for geolocation services” (P 13).

As I mentioned above, the AI Industry was partially enabled by the discovery that GPUs were a great improvement over CPUs when it comes to running AI applications. GPUs were not, however, created with AI in mind. AI producers are now investing in the development of special AI chips, in direct competition with traditional GPU manufacturers such as NVIDIA, from which they have historically purchased GPUs. Google, to take one example, has been developing the Tensor Processing Unit (TPU), which is specially designed for training and running artificial neural networks (Freund 2017). However, the “performance boost provided by TPUs works only if you use the right kind of machine-learning framework with it. And that means Google’s own TensorFlow” (Yegulalp 2017). Thus, the open source distribution of TensorFlow can be read as a technique for channeling developers towards the Google Cloud with its specialized hardware. Not to mention that the combination of both technologies gives Google control over both hardware and software for its production processes, giving the company an increased vertical integration and “a comprehensive and optimized platform to support their research and product development” (Freund 2017).

## 4.9 Labour in the AI Industry

The burgeoning AI Industry and widespread hype surrounding it have created an enormous demand for skilled AI labour. Peter Lee, Head of Research at Microsoft has described his company as engaged in a battle for deep learning talent with Facebook and Google. Lee said that Microsoft went from four full-time deep learning experts to seventy in the period 2011-2014 and would have hired more if they could (Vance 2014). The number of AI jobs in the USA grew by 4.5 times between 2013 and 2017 (AI Index 2017). Demand for AI workers continues to outpace supply. According to one study, in 2018 only 22,000 people in the world had the skills to “do serious A.I. research,” a figure over double that of the previous year (Metz 2018a). As I discuss below, this is because high level AI work requires advanced skills in mathematics as well as programming.

For the qualified, AI Industry salaries ranging from substantial to astronomical are available. According to Microsoft president Lee, in 2013: “the cost of a top, world-class deep learning expert” was similar to that of a NFL quarterback prospect (Vance 2014). However, there are a variety of roles in the AI Industry. Colson (2019) notes that “end-to-end algorithmic business capability requires many functions, and so companies usually create teams of specialists: research scientist, data engineers, machine learning engineers, causal inference scientists, and so on” (Colson 2019).<sup>45</sup> However, job titles in the AI Industry have yet to be, like the rest of the industry, distinctly defined. One data scientist I interviewed explained that “these [job titles] are ... not very well defined. And people use them in many different manners” (P 16). The lines between engineer and scientist are especially unclear. However, we can usefully make at least five divisions. From most remunerated to least, these are: data/machine learning scientist<sup>46</sup>, data engineer, data analyst, service worker and data ghost worker.<sup>47</sup>

Data scientists and/or machine learning scientists occupy the apex of the AI labour hierarchy and receive most of the AI industry hype as well as the stellar wages described above. Aghabozorghi and Lin (2016) breathlessly describe the data scientist as the “alchemist of the 21st century: someone who can turn raw data into purified insights”. A survey conducted by Stackoverflow (2018), which garnered 57, 138 responses, puts the median salary for a data scientist or machine learning specialist in the USA at \$102,000 USD, while the global median salary for the same position is \$60,000 USD. According to Indeed.com (2019a), data scientists earn an average of \$120,301 USD yearly. With the average wage in the USA in 2017 at \$50,620, data scientists are doing well for themselves.

---

<sup>45</sup> Theuwissen (2015) offers another typical breakdown: data scientist, data analyst, data architect, data engineer, statistician, database administrator, business analyst, data and analytics manager.

<sup>46</sup> Data scientist and ML scientist are often used interchangeably but are not necessarily so. According to one 240 interviewee study 90% of data scientists are involved with ML to some degree. For 40% it constitutes the majority of their work (Theuwissen 2015, 11).

<sup>47</sup> This refers only to the technical side of AI businesses. There are, of course, numerous roles on the line of business side which may be more or less involved in technical matters, and some workers do straddle both sides.

If they can manage to be hired by one of the tech giants, data scientists can do very well indeed. Someone “proficient in deep learning can earn upward of \$250,000 a year at places like Google and Facebook, according to several sources; exceptional or more experienced ones can net seven-figure salaries” (Bergen and Wagner 2015). Another account puts the figure even higher: “Typical A.I. specialists, including both Ph.D.s fresh out of school and people with less education and just a few years of experience, can be paid from \$300,000 to \$500,000 a year or more in salary and company stock” (Metz 2017a). For well-known names, the sky is the limit. As Metz (2018a) reports, Sutskever, student of Hinton, was paid \$1.9 million USD in 2016 for his position at the then-non-profit OpenAI. These high salaries are draining experts from academia. Uber alone poached 40 researchers and scientist from Carnegie Melon University in 2015 (Metz 2017a). According to Theuwissen (2015) nearly “50% of data scientists get contacted at least once a week about a new job opportunity ... 85% get contacted at least once a month” (8).

The data scientist is the one who creates machine learning models and algorithms and performs advanced analytics with them. According to Anderson (2018), data scientists often come from backgrounds in math, statistics or physics, tend to have graduate degrees, and also know how to program. Data scientists are distinguished by their central technology: machine learning, which is reflected in their kind of work. One data scientists told me “I’m expected to do more modelling and more delivering ... ML models as opposed to analytics. That is a data scientist’s job ... data comes in, data engineers process them, and then deliver them ... to the data scientists” (P 16). Another data scientist told me his work comprised three domains: “scheduled incremental feature work ... build[ing] in more data science capabilities and feature stuff that we decided on ... [and] ... firedrills. Where something isn’t working or something crashed and we need eyes on it immediately ... And then a lot of the time is also spent just for open ended research” (P 15).

Data engineers are similarly highly educated to data scientists but tend to have stronger programming backgrounds rather than maths or physics (Anderson 2018). They specialize “in creating software solutions around big data” and building “data pipeline[s]”

by connecting diverse big data technologies (Anderson 2018). They are often characterized as preparing the big data infrastructure necessary for data science (Aghabozorghi and Lin 2016). Despite not partaking in the glamour of being called a scientist, according to Indeed.com (2019b), data engineers seem to be paid similarly at an average \$129,653 USD.

Quarterback-equivalent salaries are not available to all members of the AI Industry workforce. A study from the University of California in Santa Cruz shows that “[n]ine in 10 workers in Silicon Valley make less now than they did in 1997 after adjusting for inflation” with only the highest earners seeing increases (Sheng 2018). Less lucrative data work abounds in the AI Industry. A prominent example of this are data analysts, who “query and process data, provide reports, summarize and visualize data” (Aghabozorghi and Lin 2016). They work with pre-existing software and tools to do things with data, including create visualizations, and do not generally create software. They have less education than data scientist or engineers, usually possessing a bachelors and/or professional certificates in fields such as data mining. According to Indeed.com (2019c), data analysts earn approximately half the wages of data scientists and engineers at \$65,502 USD yearly.

Finally, we have a category not always included in discussion of the AI industry workforce: data ghost workers.<sup>48</sup> The term “ghost work” comes from Gray and Suri (2019) who define it as the “often intentionally hidden” human work “powering many mobile phone apps, websites, and artificial intelligence systems” (4). Since Roberts’ (2014) exploration of the content moderation work done behind the scenes of platforms like YouTube, ghost work has become increasingly visible. Ghost workers fill in the gaps

---

<sup>48</sup> Another often invisible part of the AI Industry workforce are the service workers who provide the perks for Silicon Valley-style tech workplaces. These include shuttle bus drivers, cooks and janitorial staff. According to a report by Working Partnerships USA (2016), in Silicon Valley these workers earn a yearly average of \$19,900 USD while local rent averages \$21,444 USD (1). Most tech service workers are employed by subcontracting agencies that offer minimal job security and benefits, if any. Despite, or because of, this they are the most organized segment of labour in the AI Industry and have won some gains through unionization.

that automated systems cannot manage, identifying explicit images and identifying hate speech, to take only two instances. Ghost workers are usually hired through automated systems like Amazon's Mechanical Turk to perform small piece work tasks for micropayments – often averaging less than \$5 USD an hour (Hitlin 2016). The low pay and precarious nature of ghost work make it difficult to learn a living on. In addition, some ghost workers, like content moderators, are exposed to traumatic content which may take a psychological toll.

Ghost work is essential to contemporary AI. Singer (2019) calls it the “secret sauce for most of the AI development today”. Machine learning requires massive amounts of training data, but data must be prepared before it can be used for training. Many, though not all, ghost workers come from poorer countries where the low piece work wages can provide some degree of subsistence. The work is usually very repetitious and involves long, dedicated hours of work if it is to be a primary source of income. Yu (2017) compares this type of work to that done in Fordist factories: “today's AI projects are being built to a large degree using old-fashioned manual labour”. Nakashima (2018) describes it as “the digital equivalent of needlework —drawing boxes around cars in street photos, tagging images, and transcribing snatches of speech that computers can't quite make out”. It is, regardless of the comparisons drawn, far from the revolutionarily new type of immaterial labour posited by *post-operaismo*.

One prominent product of AI ghost work is the massively influential image dataset ImageNet, which is widely employed in training image recognition machine learning models. ImageNet is a set of 14,197,122 images with content varying from amphibians to geological formations to people. All of these images are labelled according to these categories as well as a plethora of subcategories; the category “animal” alone has 3822 subcategories (ImageNet 2010.). ImageNet was labelled by 49,000 Mechanical Turk ghost workers, hired by machine learning researchers between 2007-2010 (Li 2017). Today, professional data labelling companies exist and some are even attempting to describe their product as “fair trade” data because they offer more traditional employment structures and benefits than do platforms like Mechanical Turk (Kaye 2016). However, the title of “fair trade” data is merely self-appointed by companies looking to appear



ethical to customers and there are few prospects of ghost workers unionizing in the foreseeable future (Kaye 2019).

Even if the conditions of ghost work are improved, these workers still face a peculiar form of precarity. As they help train AI systems, these workers render their own functions obsolete and must constantly shift to new tasks. Gray and Suri (2019) hold that “[o]nce they have successfully trained artificial intelligence to perform like humans, workers move on to the next tasks engineers assign them that push the boundaries of automation” (17). This “paradox of automation’s last mile” gives Gray and Suri (2019) a measure of hope for ghost workers because, even if the work is unpleasant, “the desire to *eliminate* human labor always *generates* new tasks for humans” (17, emphasis original). According to their analysis, ghost work thus epitomizes how the threat of automation is often oversold and how a human component to AI production remains insoluble. Pointing out the essential and often hidden human labour behind AI is important for any critical evaluation. Sadowski (2018) provides several examples of this – which he calls “Potemkin AI” – not only in preparing training data, but also in the deployment of diverse AI products. To take one example, in 2017 the company Expensify, which purports to automate business document processing with machine learning, was revealed to be outsourcing the work Mechanical Turk ghost workers. These documents included receipts and benefits documents with personal information and home addresses (Gallagher 2017). While such occurrences reveal significant holes in the rhetoric of AI producers, as the next chapter shows, the long-term necessity of human ghost work is dubious.

## 4.10 Composition of AI Industry Labour

The AI Industry workforce is characterized by a heightening of the dynamics of inequality present in the larger tech sector. Further, due to the contemporary AI Industry’s focus on machine learning, which relies on training data, this inequality is directly manifest in the products of the industry.

Compared to private industry at large, the tech sector is more male and more white. While women make up 48% of private industry at large, in tech they represent only 36%.

While tech employs 6.2% more Asians than private industry in general, it employs 7% less African Americans (US Equal Employment Opportunity Commission 2016). In the AI Industry, these same dynamics are exacerbated. A World Economic Forum report shows that only 22% of “AI professionals” in the world are female (Duke 2018). In addition, female AI workers work predominantly in the use and application of AI, rather than in its development: “female AI professionals are more likely to work in ‘traditionally female’ industries – those which already have a relatively high share of female workers, such as the nonprofit, healthcare and education sectors” (Duke 2018). In addition, female AI and tech workers get paid less for doing the same work as their male counterparts. A 2016 US Labor Department investigation found that Google had “systemic compensation disparities against women pretty much across the entire workforce” (Kolhatkar 2017). While Google has disputed this, and refused to disclose all of its employees’ earnings, an employee-led investigation (only totaling 2% of employees) revealed that men do indeed make more than women (Ehrenkranz 2017). New legislation in the USA may soon force Google, and all companies with over 100 employees, to disclose wages in relation to employee gender and race (Smith, Greenfield and Green 2019).

A report from the AI Now Institute notes that only 4% of Facebook and Microsoft’s workforces are black and a mere 2.5% at Google (West, Whittaker and Crawford 2019, 3). Working Partnerships USA (2016) shows that while 10% of the “direct tech” workforce is black or Latino, the same groups comprises 58% of “blue-collar potential contract workers” (3). The poorly paid ghost workers which ground AI production are less likely to be white males. A study from the International Labour Organization shows that ghost workers – which it calls crowd workers – come from “[n]early all regions of the world ... with important representation from workers in Brazil, India, Indonesia, Nigeria and the United States, as well as Western and Eastern Europe” (Berg et al. 2018, 31). Of these workers, one in three are women, though in developing countries this drops to only one in five (Berg et al. 2018, xvi). The divide between highly-remunerated, predominantly white or Asian, male data scientists and engineers and ghost workers, who are largely non-white and often not male, is sharp.

Further, attempts to increase diversity in the AI Industry have met with substantial criticism. A recent report from the think tank AI Now judges that these attempts have been ineffective due to a narrow focus on white women, as well as focus on helping women get from school to industry, rather than addressing the sexist and racist power dynamics which characterize working in the industry (West, Whittaker and Crawford 2019, 3).

The tech industry is distinguished by a primarily white-male frat boy culture. Chang (2019) describes Silicon Valley as a “Brotopia” defined by systemic denigration of women. Chang (2018) describes the lavish parties held at the mansions of tech entrepreneurs as bastions of traditional male chauvinism in which women are treated as accessories: “outside of the new types of drugs, these stories might have come out of the Playboy Mansion circa 1972”. The widespread sexism of the tech industry became visible to the public with the publication of the study “Elephant in the Valley” which showed that over 200 female respondents experience a wide variety of discrimination in tech work (Vassallo et al. 2015). 2017 saw the circulation of Google engineer James Danmore’s internal memo that “questioned the company’s diversity efforts and argued that the low number of women in technical positions was a result of biological differences instead of discrimination” (Wakabayashi 2017). Danmore was fired, but the same year saw a massive amount of disclosures of sexual assault and harassment by female workers in the tech industry (Benner 2017a; Benner 2017b). After an outpouring of apologies by various male tech industry figures, including many CEOs, pervasive discrimination continues. A recent study by recruiting company CWJobs shows that nearly 30% of women in tech jobs have been told they only obtained their job because of their gender, while 51% say that someone has implied that their gender might prohibit their career (J. Forbes 2019).

Due to the nature of machine learning, the consequences of discrimination in tech work extend beyond the workplace experiences of female workers. As Crawford (2016) asserts, AI’s “white guy problem” is “fundamentally a data problem”. Machine learning systems “pick up any tendencies that already exist in the data they train on” (Dickson 2018). When they are produced in work environments characterized by rampant

discrimination, “[s]exism, racism and other forms of discrimination are ... built into the machine-learning algorithms that underlie the technology behind many ‘intelligent’ systems that shape how we are categorized and advertised to” (Crawford 2016).

Therefore, AI “may already be exacerbating inequality in the workplace, at home and in our legal and judicial systems” (Crawford 2016). Three years after Crawford’s piece was published, examples of perniciously biased AI systems abound; a Google algorithm categorizes black people as gorillas, predictive policing algorithms target black neighbourhoods disproportionately and a beauty contest judged by AI overwhelmingly chose white, and a few Asian, winners (Dickson 2018). An Amazon AI tool for hiring employees was scrapped before deployment when it was found to prefer the resumes of men over those of equally qualified women (Cook 2018). These are only a handful of possible examples.

One might expect that such palpable inequality and discriminatory power dynamics would catalyze a highly organized labour force, but this is not the case in the AI Industry, nor has it been the case historically in the tech industry at large. There are, however, some signs of change. Discrimination is in fact one issue around which AI labour is rallying.

## 4.11 AI Industry Labour Organization

The AI Industry has no formal labour organization among its high-skill employees. This is another trait it inherits from the tech industry at large, and from Silicon Valley in particular. Hyde (2003) notes, “[i]n high technology, unions are hardly present at all” (155). There are several factors contributing to this state. One popularly cited reason for the lack of labour organization in high-tech work are the high wages and general job satisfaction experienced by these workers (Milton 2003, 32). It is true that AI giants such as Google famously distribute free perks to keep employee spirits high, including “free meals at more than 30 cafes ... nap pods, [and] a concierge service” (ABC 7 News 2018). Another factor contributing to the lack of unionization is the high demand for, and low supply of, skilled workers. This imbalance means that it is “at least as easy for a tech worker to quit their job and find another better one as it is to attempt to organize their coworkers” (Patel 2017). While these are both no doubt significant factors, they must be

considered alongside a third factor which goes back to Silicon Valley's roots. In the 1960s, Bob Noyce, co-founder of Intel stated:

Remaining non-union is an essential for survival for most of our companies. If we had the work rules that unionized companies have, we'd all go out of business. This is a very high priority for management here. We have to retain flexibility in operating our companies. The great hope for our nation is to avoid those deep, deep divisions between workers and management that can paralyze action (quoted in Rogers and Larsen 1984, 191 in Hyde 2003, 155).

This view of labour organizing persisted throughout the tech boom and remains with the AI Industry, even if it is now most often camouflaged in what Leonard (2014) calls "stealth libertarianism". Whatever the exact conjunction of reasons, it is a fact that the only successful organization attempts in Silicon Valley have been amongst contract service workers (Hyde 2003, 155). From 2014-2017 around 5000 contract services workers became unionized in Silicon Valley. Labour organization may now be spreading to higher skill tech workers.

Political orientation is not what has kept tech workers away from labour organizing – many are left-leaning. As Roose (2013) notes, "despite being largely socially progressive and voting overwhelmingly for Democrats in national elections, Silicon Valley is probably America's least-organized labor industry". However, since the election of Trump, high-tech workers have begun to get visibly agitated, with some attending anti-Trump rallies, some protesting AI company Palantir's involvement in Trump's border policing (Buhr 2017). Trump and immigration are not the only issues which tech workers are mobilizing around. USA Google employees have also effectively protested their employer's military contracts with the Pentagon, causing the company to drop contracts and global walkout by 20,000 Google employees in 2018 after the company gave a \$90 million severance package to Andy Rubin, creator of the Android operating system, when he was fired for sexual assault on an employee (Harwell 2018; Canon 2018).

This activism did not arise spontaneously. Much of it should be attributed to the efforts of grassroots organizations such as the New York City Democratic Socialists Tech Action Working Group, Tech Workers Coalition, and Tech Solidarity. Such groups have organized demonstrations, education sessions and are aiming to influence policy (Coren 2017). They have also advised and aided tech workers in the rare attempts at unionization. Perhaps the first instance of this occurred in 2018 when 14 software engineers at the software startup Lanetix (since rebranded as Winmore) attempted to join the union NewsGuild–Communications Workers of America (CWA). Shortly after receiving notice, Lanetix management fired all of the involved employees. The retaliatory nature of the firings drew heavy media attention and support from the Tech Workers Coalition, who joined fired employees at a protest outside of Lanetix offices in San Francisco in March 2018. The National Labour Relations Board ruled in the workers' favor and Lanetix paid them a \$775,000 USD settlement (Perry 2018). This sets a precedent that other tech workers, such as those in the AI Industry, might follow. Another event occurred in February 2019, when web development tool company NPM laid off five workers who were talking to organization including the International Federation of Professional and Technical Engineers and Tech Workers Coalition about forming a union to fight degrading work conditions (Conger and Scheiber 2019). Once the issue was brought to the National Labour Relations Board, NPM paid a settlement, but no union was formed. While it is, of course, impossible to predict the future of tech work organization, at least some commentators decry a “budding socialist movement” in Silicon Valley (Spencer and Karlis 2019).

## 4.12 Conclusion

This chapter has surveyed the political economic dynamics of the contemporary AI Industry; its highly concentrated oligopolistic capital side as well as its dramatically stratified labour force, riven by rampant sexism and racism. It is in this volatile context that AI is actually produced. To complete the picture of AI production the next chapter shifts perspective to a labour process analysis.

## Chapter 5

### 5 The Production of Machine Learning: A Labour Process Analysis

The previous two chapters charted how AI evolved from a fringe research interest for a handful of scientists to a central, and expanding, industry of cybernetic capitalism. In this chapter, I shift from macro scale analysis to a micro scale analysis of the labour process by which machine learning (ML) AI is produced. The goal is to develop a concrete understanding of what “AI work” looks like.<sup>49</sup> This analysis is based on data collected during interviews I conducted with workers and management in the AI Industry between July 2017 and January 2018 and was fleshed out with follow up research in 2018 and 2019.

First, I describe my interview methodology. Second, I describe the three technical stages of the machine learning labour process. Third, I discuss four key themes drawn from my interviews. All of these relate, in different ways, to my central research question concerning the purported new autonomy of immaterial labour posited by *post-operaismo*. They are the commodity form of machine learning and its effects on AI work, the empirical mode of control of the machine learning labour process, the conception of AI as automation and the automation of AI work itself. Finally, I explore in detail the emerging method of automating AI work known as automatic machine learning (AutoML). This chapter thus sets the scene for the next in which I argue that AI work does not exhibit the new autonomy posited by *post-operaismo*.

---

<sup>49</sup> In this chapter, AI and ML may be taken as equivalent. While I distinguished earlier between ML and other types of AI, this chapter only deals with people working at companies which produce ML products. My interviewees, like the rest of the AI Industry, tended in conversation to interchange various terms including AI, ML, machine intelligence and data science. Some of them, as this chapter shows, even remarked on the fluidity with which such labels are used in the industry.

## 5.1 Group Selection

My interviewees were drawn from the AI Industry. The AI Industry comprises, as Chapter 4 discusses in detail, a handful of powerful AI tech giants such as Google, older conglomerates that are shifting gears into AI production like Siemens, a wealth of startup companies and a variety of think tanks. I aimed to recruit interviewees from each of these four categories so as to represent a cross-section of the industry. I did not, however, manage to recruit anyone from a think tank.

I recruited interviewees by contacting AI companies directly. By reading technology, business and investment journalism, I was able to compile a list of AI companies. I initially focused on Canadian AI companies, but in the course of recruitment, I chose to expand my focus worldwide. This was motivated by a lack of responses from the relatively limited supply of Canadian companies, as the following section discusses. One substantial resource I came across while expanding my list of potential companies was the database of AI companies compiled by the AI business research company Emerj (nd.).

The AI Industry is, like any other industry, composed of two groups or classes of people: labour (workers) and capital (management). While I sought to interview both classes, I was more interested in recruiting workers. This was because management's point of view already finds expression in mainstream and business journalism, while that of workers often remains unheard.

## 5.2 Recruitment

Prior to recruitment, the study (Ethics File #109130) was approved by The Office of Human Research Ethics at Western University (Appendix 4).

I emailed the AI companies collected in my list at email addresses made publically available on their websites. Some companies listed addresses for the company in general, some were for the human resources department, and some only posted the address for particular members of the executive (i.e. CEO). Whenever possible, I chose to contact the human resources department. My initial email consisted of an attached recruitment poster



along with a short explanatory blurb in the body of the email (Appendices 1 and 2). In the blurb, I briefly explained the point of the interview and asked whether the company would be willing to forward a call for interviewees to employees via email or post a recruitment poster in the workplace. I offered to send the official Letter of Information and Consent if they were interested (Appendix 3).

In all, I sent 92 emails. Of these, I received less than 20 responses. Some of these asked for more information, but were ultimately not interested in participating. However, in addition to the positive responses to my emails, I gathered a number of participants through word of mouth. Some participants and some contacts in my network shared information about my study with their contacts. This led to people volunteering to participate. When a potential interviewee expressed interest, I sent them the official Letter of Information and Consent. After the potential interviewee had reviewed this document, I obtained verbal consent from them before starting the interview. Since the interviews pertained to work, and there could be negative consequences for employees expressing critical views of their employers or workplaces, all interviews were conducted under conditions of anonymity.

### 5.3 Interview Preparation

I employed semi-structured, qualitative in-depth interviews. The qualitative interview is a “knowledge-producing conversation” through which we can learn about “how [people] experience the world, how they think, act, feel and develop as individuals and in groups” (Brinkman 2013, 1). I chose the qualitative method because I wanted to learn about people’s experiences working in the AI industry and not just quantitative facts.

I adopted a semi-structured approach because I wanted to remain open to change in the course of the interview. In semi-structured interviews, the interviewer produces a loose plan or series of topics which are intended to guide the interview, but rigorous adherence is not sought. Instead, the interview is conducted as a conversation and is allowed to develop organically, within certain bounds (Brinkman 2013, 21). The goal of a conversational attitude is also reflected in the qualifier “in-depth,” which signifies a closely engaged, relatively informal interview format intended to establish rapport

(Johnson and Rowlands 2012, 99). I chose this approach because it is “best suited to research questions of the descriptive or exploratory type” (Johnson and Rowlands 2012, 101). Since very little information was available about the AI Industry, exploration was precisely what I wanted to do. Further, in-depth interviews “rarely constitute the sole source of data in research. More commonly, they are used in conjunction with data gathered through [other] avenues” (Johnson and Rowlands 2012, 100). In particular, they may be used as “a way to check out theories” (Johnson and Rowlands 2012, 100).<sup>50</sup> Assessing the validity of immaterial labour theory was precisely my goal.

I prepared a list of interview topics grouped under four broad categories: labour process/organization, industry dynamics, use of machine learning to produce machine learning and speculations. The list originally consisted of ~30 questions, but after the first few interviews, I pared this down to a more manageable 16 (Appendix 5). Some of these eliminations were based on redundant answers, while one (regarding income) was dropped due to the refusal of the first three participants to answer.

## 5.4 Conducting the Interviews and Composition of Interviewees

Ultimately, I conducted 16 interviews between July 2017 and January 2018. One interview was discarded. That interview was with a manager at a regional utilities company that had recently begun experimenting with AI in their operations. While it was interesting, it did not meet my criteria for belonging to the AI Industry since that company did not produce an AI product.

Thus my research draws on a total of 15 interviews. The interviews were an average of 60 minutes long. One interview was conducted at a café on a university campus while the 14 other interviews were conducted via Skype, Google Hangouts or telephone and were recorded with a digital voice recorder and later transcribed and anonymized by myself.

---

<sup>50</sup> I chose interviews instead of surveys because of these exploratory capacities. While focus groups could have been useful, the expected (and confirmed) difficulty in recruiting AI workers and management for one-on-one interviews would have been substantially compounded by having to coordinate multiple participants’ schedules and/or physical locations.

Brinkman (2013) notes that 15 is a common amount of participants for interview studies as it presents a manageable quantity of data and since “the aim is not statistical representativeness ... but instead the chance to look in detail at how selected people experience the world” (59). However, it is important to note that if samples are too small “they can easily miss key constituencies within the population, or contain too little diversity to explore the varying influences of different factors” (Ritchie, Lewis and Elam 2003, 85). Or in other words, “small-scale samples only work in qualitative research if good purposive or theoretical sampling has taken place” (Ritchie, Lewis and Elam 2003, 84-85). This was precisely why I sought to interview both workers and management.

Capital or management was represented by 5 participants, all of which were CEOs. Labour was represented by the remaining 10 participants, who ranged across the hierarchy of AI Industry labour roles. 4 participants defined their roles as scientist (2 data scientists, 1 machine learning scientist, 1 lead scientist). 3 participants defined their roles as engineer (2 senior software engineers and 1 inference engineer. 1 participant identified as a research and development programmer, 1 as an intern and 1 as a PhD student at a research institute with close ties to industry. Ghost workers are, unfortunately, not represented in these interviews. This is because, at the time of conducting the interviews, I was unaware how critical such work is to the production of AI. Thus, these interviews focus on the high end of the spectrum of work in the AI Industry and are thus, as the previous chapter discusses, drawn from an overwhelmingly white and male pool of possible participants.

I did not ask interviewees to identify their race, but interviewees were residents of North America, Africa and Asia. The participants ranged in age from 27-43 with an average age of 28.6 years. All participants were male. Despite an effort to reach specifically female participants, including directly e-mailing several prominent female public figures in AI, I did not obtain any. I hoped to speak to female AI workers to gain their perspectives on the substantial gender bias exhibited by the tech industry, as discussed in Chapter 4.

13 participants had obtained or were in the process of completing a Master’s degree, while 6 had obtained or were in the process of completing doctorates. Almost all of the

participants had studied or were studying a type of science or engineering, but particular fields varied from cognitive science, to computer science, to mathematics, to quantum mechanics, atmospheric science and medical science.

8 participants worked at startup companies. Half of these had less than 10 employees and half had between 10 to 40 employees. 3 worked at startups which were either distinct entities created within a larger firm or were bought up and integrated into a larger firm. These internal startups had between 10 and 19 people. 2 participants worked for large international companies. The particular branches they worked at employed between 30-50 people. 1 participant was at a research institute with around 20 people and 1 worked at a small but established company of around 10 people. The majority of these companies were established between 2012-2017 with only the 2 large international firms dating to 1998.

Most participants had not been at their current position for a long time. Only 2 had held their current position for 5 years. 4 had been at their current position for 2-3 years, while 9 had been there for around a year or less.

## 5.5 Exploratory Interviews and Follow-up Research

As mentioned above, the interviews were conducted in an exploratory mode. As an outsider to the AI Industry and data science generally, I learned about a lot of new things in the course of the interviews. After the interviews, I conducted follow-up research regarding the technologies, practices and other phenomena that I heard about. The presentation of my interview data in this dissertation is supplemented by, and integrated with, this follow-up research.

## 5.6 The Machine Learning Labour Process

In Chapter 2 I discussed how in a capitalist economy a labour process is “in general a process for creating useful values” which is also “specifically a process for the expansion of capital, the creation of a profit” (Braverman 1998, 36). While AI is often portrayed as an abstract, intellectual enterprise, it is, like any other commodity, produced by labour. As the AI Industry expands, recognition of this is growing. In a recent business-oriented

lecture, Oren Etzioni (2018), computer scientist and serial AI entrepreneur emphasizes that making AI requires “99 percent manual labor”. Thompson (2019) notes that while the stereotypical image of a software producer has been the “hoodied young Zuckerbergian,” a new image is emerging – that of the “Blue-collar Coder” (326; see also Dash 2012). The rest of this chapter details the labour of these in the AI Industry. Any labour process is composed of individual days of work. When asked to describe a typical day at work in the AI Industry, one R&D programmer working at a startup told me:

You usually have some pending tasks from a previous project, or you have a new task from a new project ... the CEO comes over to your desk and he starts talking about what exactly the vision is and what the project is about. How we should do things. He discusses certain algorithms and various concepts that need to be ... fleshed out before it can be done in the code. There can be multiple people standing around a computer and talking about various ... high level design principles. Then they can go to a white board and we can talk about ... various little details ... [but] [m]ost programming in any company is a solitary activity ... that’s the nature of programming (P 4).

My interviews revealed a consensus on this basic type of workday breakdown – a coordinating meeting opening the day, followed by largely solitary work, interspersed with meetings. Machine learning projects are not completed in a day, however. When asked to describe a typical day at work, one machine learning scientist told me that “Due to agile development ... it makes more sense to talk about [work] in terms of the week. That is, our sprints” (P 16). Agile is a software development methodology which nearly all of my respondents described their workplaces as operating, to a more or less formal degree. Sprints are periods between a week and a month long into which development is broken down. I discuss Agile and other development methodologies later, but first I detail the technical stages of the machine learning labour process.

The machine learning labour process is typically broken down into three stages. Dettmers (2015) writes that with machine learning: “we (1) take some data, (2) train a model on

that data, and (3) use the trained model to make predictions on new data”. Similarly, Dong (2017) explains that, “[m]achine learning engineering happens in three stages — data processing, model building and deployment and monitoring” with model building being the “meat” of the machine learning sandwich because this is where predictive capabilities emerge (Dong 2017). However, the bread actually takes more labour time to produce. According to one study, data scientists spend 19% of their time at work collecting data sets, 60% cleaning and organizing it and merely 7% building training sets and refining algorithms (Crowdfunder 2016, 6).<sup>51</sup> Let us now run through the three stages in detail.

## 5.7 Data Processing

Data processing “involves cleaning and formatting vast amounts of data to be fed into the model” (Dong 2017). This is the primary stage for ghost work. Ready-to-use data does not exist and a lot of work must be done to make data usable. Data is often extracted from databases with tools such as SQL. Brownlee (2013) breaks this stage down into three sub-steps which he describes as “very likely to be iterative with many loops” (Brownlee 2013). First is selecting data or determining “what data is available, what data is missing and what data can be removed”. While more data is generally better for machine learning, not all available data are going to be relevant to every problem. Additionally, important data may be missing and may have to be simulated. Second is preprocessing data “by formatting, cleaning and sampling from it.” This involves making the data usable by the tools that process it, fixing errors or missing values and removing sensitive information. In addition, huge datasets may need to be sampled rather than used whole to reduce the computational load (Brownlee 2013). Dong (2017) describes this stage as “frustrating manual labor” and “repetitive work”. It is, however, necessary. As one R&D programmer told me, “[i]f you want to get very good results then you need to clean out the data a lot” (P 4). Third, comes feature engineering or data transformation: working up preprocessed

---

<sup>51</sup> In addition to the direct labour process, ML work involves a lot of research. According to Theuwissen (2015), 83% of interviewed data scientists “spend between 25-75% of their time on R&D as opposed to production” (12).

data by engineering features using techniques including “scaling, attribute decomposition and attribute aggregation” (Brownlee 2013). Feature engineering is “the process of transforming raw data into features that better represent the underlying problem to the predictive models, resulting in improved model accuracy on unseen data” (Brownlee 2014). This is an essential step because “[m]uch of the success of machine learning is actually success in engineering features that a learner can understand” (Locklin 2014). For Dettmers (2015) it is the “most important skill” for making machine learning. Because not all features are obvious, feature engineering is often referred to as an “art” requiring some sense of intuition (see also Brownlee 2014).

## 5.8 Model Building

In this stage the processed data is input to a learning algorithm which then produces a model which “contains the learned relationships” between the data (Brownlee 2015). This model can then be applied to analyze unseen data. Tools used here include TensorFlow and Spark ML. Here it is useful to distinguish between three primary types of learning: supervised, reinforcement and unsupervised learning.

The majority of commercial machine learning today uses supervised learning (Brownlee 2016). In supervised learning, the data from which the algorithm learns is labelled by humans, usually in terms of categories. This approach thus “require[s] a feedback signal (from external sources) ... in order to get going” (Sun 2014, 111). The system learns the categories by discerning patterns across the labelled examples. The labelled data thus acts as a supervisor (Brownlee 2016). Given enough labelled photos of red hexagonal signs with the word STOP on them, in various visibility conditions and from various angles, as well as examples of non-stop signs, a supervised learning system can learn a concept of stop sign and output the category ‘stop sign’ whenever it is fed an image of a stop sign. Ideally, the concept of stop sign formed by the system will be robust enough to enable it to generalize beyond its training data, to recognize stop signs in images it has never seen before. The “supervisor provides the correct values, and the parameters of a model are updated so that its output gets as close as possible to these desired outputs” (Alpaydin

2016, 111). Learning ceases when “the algorithm achieves an acceptable level of performance” (Brownlee 2016).

Because of the necessity of labelling, supervised learning entails a lot of repetitive labour (often performed by ghost workers, as discussed in Chapter 4). AI companies have, however, developed creative methods to obtain labels without paying even miniscule ghost work wages. The reCAPTCHA system (acquired by Google in 2009), which is designed to prevent bots from accessing certain websites, asks human users to label objects such as stop signs, vehicles and storefronts in photos. The click work involved is, by Google’s (nd) admission, used to “build machine learning datasets” and help “solve hard AI problems”. Some machine learning practitioners speak of a “machine learning data bottleneck” and lament that “[l]abeling data by hand can be time consuming, expensive, and impractical; and sometimes you don’t even have sufficient examples to label, especially of the rare events that are most important” (Mugan n.d.). The two other types of machine learning may be seen, in part, as attempts to overcome this data bottleneck.

In unsupervised learning “there is no predefined output, and hence no supervisor ... only the input data” (Alpaydin 2016, 111). Instead the goal is to automatically “find structure in the data” (Alpaydin 2016, 117). In so doing, unsupervised learning may be said to generate categories or theories – such as that of stop sign. In other words, unsupervised learning automatically produces a “hidden model” which represents a theory of the meaning of data, or its “underlying factors and their interaction” (Alpaydin 2016, xi).<sup>52</sup> As such, unsupervised learning can be considered the automation of supervised learning. Indeed, the work entailed by labelling data drives research in unsupervised learning because “unlabeled data is a lot easier and cheaper to find” (Alpaydin 2016, 117). The internet is a gigantic trove of unlabeled data.

---

<sup>52</sup> Two prominent types of unsupervised learning are “clustering”, in which data points are grouped by some type of similarity and “association”, in which “rules” that describe correlations between data points can be discovered (Brownlee 2016). Clustering may help support a theory about relationships between variables, but more interestingly, “there may be a cluster or clusters that no expert could have foreseen” (Alpaydin 2016, 115). Novel knowledge may be produced.



The pioneers of deep learning argue that unsupervised learning will eventually become the central machine learning approach because it is how humans and animals evolved to learn; not by being told what everything in the world is called, but through observing it (LeCun, Bengio and Hinton 2015, 442). One group of researchers recently pointed out that high definition video contains so much data that to analyze it with supervised learning, with its need for labelled data, would require “several orders of magnitude more labels” than doing the same for pictures, and as such seems infeasible (Luo, Peng, Huang, Alahi and Li Fei-Fei 2017, 2203). They suggest the unsupervised learning will become a necessity for AI.

However, unsupervised learning finds competition in the approach called reinforcement learning, which attempts to overcome the data bottleneck in a different way. The pioneers of this approach describe it as “learning what to do ... so as to maximize a numerical reward signal. The learner is not told which actions to take ... but instead must discover which actions yield the most reward by trying them” (Sutton and Barto 1998, 127). Alpaydin (2016) describes reinforcement learning as “learning with a critic” (127). However, the so-called critic is not human. Instead, reinforcement learning aims to emulate the process of learning by experience, just as young humans learn that fire is painful to the touch through trial and error. There “is no external process that provides the training data. It is the agent that actively generates data by trying out actions in the environment and receiving feedback (or not) in the form of a reward” (Alpaydin 2016, 128).<sup>53</sup>

Reinforcement learning was thought only be usable in simple domains, such as Backgammon, which it mastered in 1992, until 2013 when DeepMind combined it with unsupervised learning (Knight 2017). DeepMind showed that this combination of approaches could be used to teach a system to play Atari games with superhuman skill,

---

<sup>53</sup> The necessity of experimentation means that reinforcement learning may require less data but more time. However, because such systems can operate in virtual environments at accelerated timescales, they can accrue much more experience (data) than humans can in a given period. Indeed, Andrew Ng holds that to accrue the necessary data for successful reinforcement learning such systems must “practice relentlessly in simulations” (Knight 2017).

without programming any knowledge about the games into the system and giving it access only to the pixel information displayed on the screen. The same combination enabled AlphaGo's win over Go master Lee Sedol in 2016.

## 5.9 Deployment

A trained machine learning model is deployed by integrating it “into an existing production environment in order to start using it to make practical business decisions based on data” (DataRobot n.d.). Tools such as Git, Docker and Grafana are used. One major hurdle of deployment is that there is often a discrepancy between the machine learning programming languages and other business software (DataRobot n.d.). This entails either re-coding the entire model or building an API (application programming interface), which translates between the two and allows the model to be integrated with other software (Paul 2018).

Deployment also involves maintenance. The model may need to be updated on new data to reflect changes in its environment. Therefore “engineers re-train production models on fresh data on a daily to monthly basis, depending on the application” (Dong 2017). In addition, Dong (2017) explains that since “traditional unit tests — the backbone of traditional software testing — don't really work with machine learning models, because the correct output of machine learning models isn't known beforehand”. Instead, “engineers take a less structured approach: They manually monitor dashboards and program alerts for new models” (Dong 2017). Dashboards are representations, often graphical, of various metrics of the model's functioning, often in real-time. In business settings, these will often be key performance indicators (KPIs). One data scientist described to me how dashboards have simplified deployment: “Compared to the standard, years ago, you'd have ... analysis reports ... numbers on the screen ... the evolution of that is the dashboard, which is dynamic. Any person can go spend time and find” what they need (P 9). Dashboards are also often used to present models to nontechnical interested parties including management and customers.

These are the basic technical stages of the machine learning labour process. But how is this labour process influenced by the valorization process – the necessity of increasing

capital – on which it necessarily supervenes? The next sections elaborate four themes extracted from the interviews which answer this question in different ways.

## 5.10 The Commodity Form of AI

AI Industry products are commodities and their production has been structured for speed and efficiency in accord with the competitive dynamics of capitalist production. For one startup CEO, the near term goal is to “knock down some more customers” (P 2), while for another, “the whole idea is to make money through artificial intelligence but at the same time find creative and holistic ways to give back to the society” (P 5). Yet another CEO describes his company’s goal as:

Make as much money as possible in as short a time as possible ... There’s no such thing as enough ... honestly, we’re making a magic factory. And you can quote me on that. What we’re trying to do ... is that machine that shoots out baseballs and just throws strikes all day (P 13).

The commodity nature of AI entails the long workweeks familiar to software work in general. Most of my interviewees reported working well over a typical 40 hour work week. As a senior software engineer put it, “if you want to work from 9 to 5, I mean it’s fine, but don’t expect that will be very competitive in the marketplace” (P 7). An inference engineer who returned to the company he did an internship at, summed up his experience like this: “the CEO ... said the amount of hours [40] you’re putting in were what we call like mediocre. You need to put more in if you come full time. So he just made that known and I agreed to it ... I’d say maybe 50 to 60 hours, that’s low” (P 10). One CEO told me that concerning his work week, “there’s no fixed number to it, but it’s definitely above ... 80 hours a week”. In addition to these long standard weeks, AI work is subject to sporadic periods of intense, nearly non-stop work. Such so-called “death marches” have been a noted and feared component of software production since the 1990s (Yourdon 2004). As one inference engineer told me, “when stuff broke or hit the fan or there was a deadline ... it could easily go to 12 hours every single day. Including weekends. And there had been times as well when something breaks in the middle of the night. You get paged ... You have to wake up and fix it”.

The competitive nature of machine learning production also means that work is changing rapidly as related technologies continue to evolve. This is an industry where work is subject to continual revision and workers must continually learn new skills. One machine learning scientist described his work as always in flux:

For a data scientist it's already very fast paced and evolving every second ... New articles from even the most prominent researchers are coming out twice a week. That changed the entire perspective of the field. Or new applications that change the entire perspective of the field. So it's already very fast paced. So if ... it's a data science job, it's an everyday story that the methods that you used previously are really not as good anymore ... a lot of things are unknown. We don't know how to use all these potentials in the field ... the thing that is state of the art two weeks ago is kind of taken for granted now (P 16).

One startup CEO agreed, describing the industry as having a “rate problem. It's not a fixed benchmark of what is enough, it's really how fast can you build. Because everything is changing. Nobody can predict in two years what the landscape is really going to look like” (P 2). This compels many machine learning workers to augment their university education with online courses from providers like Coursera and Udacity. One data scientist told me he was “continuously improving” his skills on such sites (P 8) while another, who reported spending around 50% of his time at work on research, said “I'm constantly trying to learn. New things are coming up. It's very difficult to keep up ... with the new developments” (P 9). Further uncertainty is added by machine learning's reliance on data which entails a whole suite of problems surrounding the quality and availability of data. According to a startup CEO: “If you're too inflexible in data science you're setting yourself up to fail. Because the truth is, you don't know what's going to happen. You hope what's going to happen. You build a plan around it that makes a lot of sense, but often times, weird things happen. For example, you think you have the data and you don't have the data” (P 13).

A common way of structuring AI production, to control the various contingencies mentioned above, is the set of development principles called continuous integration. Continuous integration means that “when someone checks his or her revised code into the repository, an automated system picks up the change, checks out the code, and runs a set of commands to verify that the change is good and didn’t break any-thing” (Meyer 2014, 14). Many of my interviewees explicitly reported using it, while it is implicit in the Agile and Agile-like methods mentioned by others. One CEO told me the general idea of continuous integration is:

you wanna be releasing as frequently as possible. So any time someone adds code to the code base that is polished and good it should go to production right away ... that should all be done automatically. People shouldn’t have to manually push code to the server. Which ... is an error prone step ... you wanna just have like an automatic flow (P 2).

An R&D programmer said that “when somebody finishes their code and pushes their code it’ll automatically get built and they’ll get the build out of that. They won’t have to build it on their machine ... It’s very essential. Without continuous integration it’s very difficult to get any work done” (P 4). One CEO reports doing continuous integration “at an unusual speed. So many companies have like a nightly build or a weekly build ... We build ... for every change somebody makes” (P 2). Continuous integration aims to automate, and thus accelerate, testing for bugs. For one commentator this is necessary because, “[i]f tests run longer than 10 minutes, developer productivity drops, slowing down the process of shipping new features or bug fixes to the customer” (Meyer 2014, 14). The commodity form of AI not only influences the labour process via competition dynamics, however. It also exerts influence through other interested parties including customers, funders and management.

The customer’s influence on the AI labour process can be substantial since many AI companies offer a consulting service as part of their business model. Others continually must respond to changing customer demands throughout the development process. One data scientist reported spending 40% of his time at work “responding to ... ad hoc

requests coming from clients” (P 9). One startup CEO told me that, although he has a technical background, he spends a lot of time:

going to conferences, going to events where we have to tell people how data science is different from BI [business intelligence] work. How data science is actually different ... from any of the software-as-a-service solutions that they are using ... You can call that business development or you can call that building up client’s knowledge base (P 6)

He elaborated that the “consulting business requires immediate response to clients, immediate photo ops, only then are you able to build a name big enough that you can ... actually get word of mouth and clients start coming back to you. Rather than you going out and pitching to clients” (P 6). A different CEO even works as a CTO for one of his largest clients as a way “to offer something extra special” (P 13). Yet another CEO emphasized the commercial aspect of his work, noting that “the two things I spend most of my time on are pitching investors and sales” (P 3). A different CEO described a downside of his startup’s consulting business model:

you know how they say that the sun never sets on the British empire? The same kind of deal. Because people are all in different time zones, it’s really bad for my sleeping patterns ... that’s a major drawback of being a high availability consultant ... it’s not healthy for you. We’re solving that with money. We got a cleaning lady who cleans the house. We’ve got a babysitter who watches the kids. So like, you know, a lot of things can be solved just with money (P 13).

Another possible vector of influence is the potential for startups to be bought up. While none of the startup employees I spoke to told me their companies were actively seeking acquisition, one CEO (who said his company was too new to consider acquisition) told me that:

one of the things that I’ve heard other companies say is that they kind of realized that they couldn’t get to the next phase on their own. So you

know the cost of acquiring new customers or they didn't have the level of skill in their management to go after some big clients. That being part of bigger company just helped them pursue their goal better ... generally there's a strong desire for companies to exit, whether that means going public or becoming much, much bigger ... being acquired (P 3).

Further influence on workers comes from management, VC funders and incubator programs, to whom machine learning workers have to explain and justify their work. One data scientist described how after the production process is complete, he has to present his work to his managers:

usually we sit in a conference room and we display the result on the projector. We zoom in to the visualizations to see what is happening. I present to higher management and after the work is approved, we compile it into a report or dashboard and then the higher management goes to the user or the customer ... you have to do a number of iterations and then come up with some result and you have to convince your upper management that this is how I did it and this is the significance of that result (P 7)

A lead scientist working at a startup embedded within a large company told me that "we report to an angel board that gives us our funding and funding decisions. But, ultimately they are less fickle than an outside agency might be" (P 15). One data scientist told me that an important part of his job was explaining how machine learning works to his bosses: "My managers don't understand the complex mathematics ... I am the one who decodes all of this for them" (P 8). One engineer told me that one of the biggest problems facing his company was that they do not have a co-founder which means that the social "bandwidth" of the company is solely provided by the CEO (P 4). A CEO described his experience in an incubator program like this:

every 8 weeks you meet a whole bunch of successful entrepreneurs, investors and they grill you. And if one or two is willing to spend 4 hours

of their own time over the next eight weeks, then you can stay in the program. If no one wants to commit time to you, then you get kicked out. So ... obviously, it can be pretty anxiety provoking and stressful (P 3).

A recent study confirms that communicating to one's superiors is a large component of work in the AI Industry. According to Bowne-Anderson (2018) "the vast majority" of data scientists say that their "key skills" are not their technical machine learning or data science skills but rather their "abilities to learn on the fly and to communicate well in order to answer business questions, explaining complex results to nontechnical stakeholders ... and convincing decision makers of their results".

In sum, the commodity form of AI necessitates a labour process optimized for speed, with long workweeks, continual adaptation of workers to technical change and an emphasis on communicating technical aspects of work to non-technical stakeholders including management.

## 5.11 Empirical Control of the Machine Learning Labour Process

The machine learning labour process is not only optimized for competitive commodity production. It has also been structured for maximum control, albeit of a decentralized or deterritorialized variety of control. This is evident if one observes the development process ideologies and the software tools employed.

As I mentioned above, nearly all of my interviewees reported their workplace as using an Agile or Agile-like development framework. Agile is a methodology for managing knowledge workers. To understand the significance of Agile, it must be seen in contrast to the waterfall development model which preceded it. The waterfall model is a sequential model of development where a certain number of stages follow one another in a rigid order. It derives from manufacturing, but was deployed in early software work before specialized paradigms for software had been developed. An early, possibly the earliest, formal description of the waterfall model appears in Royce (1970) who says it is



“risky and invites failure” in the context of software production because testing only occurs at the end, meaning that failure entails restarting the process entirely (329).

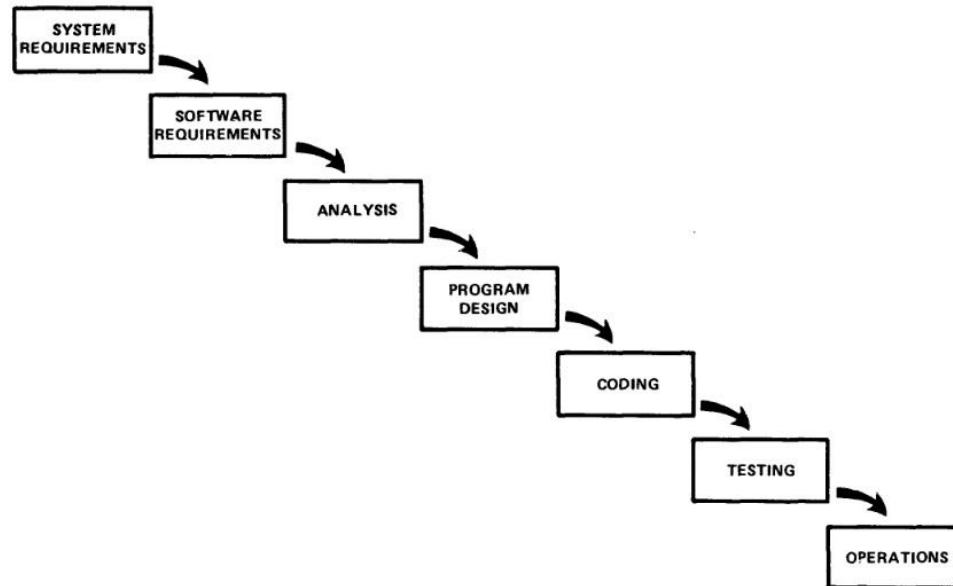


Figure 2. Implementation steps to develop a large computer program for delivery to a customer.

### Figure 1: Waterfall Methodology (Royce 1970, 329)

The problem is that a software product cannot usually be specified in sufficient detail in advance, like a bridge or automobile can be. While the waterfall method “assumes that variations are the result of errors,” Agile expects that “external environmental changes cause critical variation,” or in other words, that changes in the course of development are inevitable and their cost should be minimized (Highsmith and Cockburn 2001, 120). Proponents of the approach thus describe Agile as “the ability to create and respond to change. It is a way of dealing with, and ultimately succeeding in, an uncertain and turbulent environment” (Agile Alliance n.d.). It is thus no surprise that it has been widely adopted in today’s AI Industry with its focus on machine learning production which is “significantly ... less predictable process than traditional software development since the models learn from data rather than specific human instruction” (Yao 2019). The principles of Agile were stated in a manifesto by a group of software engineers in 2001 as follows:

Individuals and interactions over processes and tools  
 Working software over comprehensive documentation  
 Customer collaboration over contract negotiation  
 Responding to change over following a plan (Beck et al. 2001)

Rather than following a one-way waterfall sequence, Agile functions incrementally and iteratively. Incremental means that each step of the production process should produce a new working part that is added to preexisting ones. Iterative refers to recognizing the “impossibility (or at least improbability) of getting a feature right the first time” (Cohn 2010, 257). Agile development is thus organized into a series of “sprints” which are usually between a week and a month long, at the end of which the workers are expected to deliver working software (Cohn 2010, 258). The goal is that a working prototype is produced as soon as possible and is progressively improved. One software engineer summed it up to me as “the idea of pushing something out early and iterating on it” (P 12).



**Figure 2: Agile Methodology (Kuruppu 2019)**

There are many Agile approaches to software development but one that several of my interviewees mentioned (and perhaps the most popular overall) is Scrum (Denning 2015; Scrum Alliance 2018). Proponents of Scrum describe it as a “radically different approach” to software development which “reintroduces flexibility, adaptability, and productivity into systems development” (Schwaber and Beedle 2002, 1). It aims to the

increase the flexibility attributed to Agile approaches by a further dismemberment of traditional management practices. While Agile development in general may still be structured in a traditional hierarchical manner, in Scrum the workers, called the Scrum Team, are “self-organizing and fully autonomous. They are constrained only by the organization’s standards and conventions, and by the Product Backlog that they have selected” (Schwaber and Beedle 2002, 9). The Product Backlog is a “prioritized list of project requirements” which is continually modified throughout production (Schwaber and Beedle 2002, 7). The Product Backlog is dealt with through a system of commitment making. After deciding on a task and timeline, one inference engineer told me that management “treated it as a promise. Where if you missed it you’d have to give daily updates or hourly updates ... And there’d be someone else, another engineer, who’d co-sign with you. And they were also responsible for you finishing that task” (P 10).

Since the Scrum Team determines how to complete a sprint autonomously, the function of management is redefined. In Scrum, rather than direct and delegate, “[m]anagement's new and primary job [is] to maximize the team's productivity, to be there to help it do the best that it [can]” (Schwaber and Beedle 2002, 7). A primary function of a manager in Scrum is to run the Daily Scrum. This is an approximately 15 minute meeting, usually held in the morning, in which the Scrum Team assesses their progress, revises the Product Backlog, and commits to their next tasks. People not in the Scrum Team may attend these meetings, but may not speak or interfere in any way (Schwaber and Beedle 2002, 42). The manager largely reacts to what is learned from the Scrum Team during these daily meetings. Proponents of Scrum argue that by “stripping away cumbersome inappropriate and cumbersome management practices, Scrum leaves only the essence of work ... Although the Scrum process seems simple and skeletal, it provides all the necessary management and controls to focus developers and quickly build quality products” (Schwaber and Beedle 2002, 10).

The workers I spoke to generally enjoyed the autonomy provided by Scrum and Agile approaches. One research intern described his company as “very good about letting people manage their own time ... They were pretty flexible. As long as your targets are being met. And you’re making progress. I really liked that part. I didn’t have to actually

be at work. There would be days where I would just work from home” (P 17). One lead scientist expressed an interesting bipolarity of autonomy entailed by these development methods:

On the one hand, what we are doing is decided ... So on that level there is no autonomy because the projects are broadly set. But on the level of how exactly we're doing that I have complete autonomy. The absolute freedom to choose between techniques and approaches ... Or to do research into them if I feel it's necessary. That's one of the reasons I like working at a smaller place like this. I feel you get to set what the project is and you get to do what you're passionate about, what you like (P 15).

This agrees with what Barret (2005) describes with her notion of “technical autonomy” in software work, as discussed in Chapter 2 (95). However, while workers do enjoy this sense of autonomy, Scrum remains a “management and control process” and a “kind of social engineering aiming to achieve the fulfillment of all by fostering cooperation,” as proponents of it admit (Schwaber and Beedle 2002, 1). And like other management control techniques, Scrum has an interesting relationship to automation.

Scrum promoters Schwaber and Beedle (2002) distinguish between defined and complex production processes. Defined processes are “simple with unobtrusive noise” and can be controlled and executed repeatedly through rigorous definition (94). Any process that cannot be described in sufficient detail to be controlled predictably and repeatedly is a complex process that must be controlled by empirical process control. The “empirical model of process control ... provides and exercises control through frequent inspection and adaptation for processes that are imperfectly defined and generate unpredictable and unrepeatable outputs” (Schwaber and Beedle 2002, 25). They offer this illustration:

Chemical companies have advanced polymer plants that require empirical controls. Some chemical processes haven't been defined well enough for the plant to operate safely and repeatably using a defined process control model. Noise has rendered statistical controls ineffective. Frequent inspections and verification are required to successfully produce a batch.

As chemical processes become better understood and the technology improves, the plants become more automated. However, assuming predictability too soon is the recipe for an industrial catastrophe (Schwaber and Beedle 2002, 100-101)

The authors argue that because all extant software development methods are based on “partial and weak definitions of development activities” software production must likewise use empirical control (Schwaber and Beedle 2002, 95). Hence, the surveillance inherent in the Daily Scrum. Rather than dictate the work of expert machine learning workers, with Scrum management monitors them, at least until the automation of said work becomes feasible. The technical autonomy of machine learning workers must be understood in this context of control and surveillance. The words of a VP of engineering responding to a presentation on the Scrum methodology are instructive: “Fine ... I'm willing to take the risk of giving the team autonomy for defined periods” (Schwaber and Beedle 2002, 82).

Scrum can more easily be grasped as a technique of control by looking at the software tools it employs. Scrum-based production typically employs a variety of proprietary software. One of these which several of my interviewees reported using is JIRA. JIRA is a project management and bug tracking software produced by the Australian enterprise software company Atlassian. Nevogt (2019) writes that “Agile methodology is built right into JIRA”. JIRA is a centralized project management system in which “[a]ll projects are logged into a central database and each one goes through a number of workflows (processes) ... [which] control the status of the project as well as the rules by which it transitions to other statuses” (Nevogt 2019). Workflows are comprised of issues that are represented as boxes (states) connected by arrows (transitions). JIRA can automatically generate a variety of visualizations and summaries tracking various aspects of the labour process for managers. Further, JIRA includes “feature-rich time tracking” and the option for automatic screenshots, as frequently as every minute, for all users so that management can “get accurate timesheets for [the] whole team” (Nevogt 2019). Advocates of Scrum are proud that there are “no mechanisms in Scrum for tracking the amount of time that a team works. Teams are measured by meeting goals, not by how many hours they take to

meet the goal” (Schwaber and Beedle 2002, 73). JIRA, however, provides a technical way to include time tracking in Scrum and allows real-time surveillance of the work of the whole Scrum Team.

In sum, while the machine learning labour process is free from defined control models, it is subject to empirical control via development methodologies like Agile and Scrum and software such as JIRA.

## 5.12 AI is Automation

Patel (2017) writes that many “though not all, tech workers are integral to the employing class’s current efforts to flexibilize, speed up, or eliminate other people’s work”. When I asked AI Industry workers whether they thought of AI as an automation technology and whether they were concerned about its potential effects on employment in general, responses varied. One software engineer defined machine learning not as automation but rather as machines doing things that people cannot, but he was the exception to the rule (P 4). Most recognized AI as an automation technology, although there was a range of evaluations as to its likely effects on society. Nearly all respondents recognized that AI automation would have effects on employment, though the extent of effects they expected varied. What was most interesting to me was how AI-based automation was described, regardless of its potential effects. Several interviewees spoke of AI as automation in terms of desirable efficiency.<sup>54</sup> A software engineer told me:

people in the industry tend to think of ML, or AI ... in terms of efficiency ... We’re discovering ways of making processes more efficient than they currently are ... the way that technology contributes to the economy, it is in discovering new efficiencies ... Uber is cool not because people are driving around or picking people up, but because it’s more efficient than what the current system, or the preexisting system of getting rides was (P 12)

---

<sup>54</sup> This is likely connected to the notion of optimization which Wu (2019) argues is uncritically deployed widely in computer science.

The same engineer elaborated: “I definitely envision the work we’re doing enabling automation in areas which are currently busywork ... Should we pay someone 4, 5, 6 hundred thousand dollars a year just to say this is where the cancer is? That’s a cost to our health system ... In the long run, probably we shouldn’t be paying that if we can automate that” (P 12). This sentiment was echoed by a lead scientist for whom “automation itself is about, broadly, timesaving” (P 15) and by a startup CEO, who said:

It’s about making processes more efficient. Which is what software is about in general. It’s just a new kind of software. It’s something people have to adopt to stay competitive ... I sleep very well at night. I don’t worry about ... am I taking people’s jobs away? Ask your blacksmith (P 13)

One lead scientist said that fears of technological unemployment are “always rooted in the lump of labour type fallacy where this certain demographic is going to be completely replaced and therefore those people will just lose their jobs ... I think historically that hasn’t really panned out” (P 15). A software engineer, discussing automation, said: “I am an optimist. I think if there are no sweatshops left, that would on the whole be a net good for humanity” (P 12). An inference engineer similarly told me that he expects the “overall impact in society ... will be positive in the long run ... I think it has potential to free up people’s time so they can maybe do something else that’s productive for the economy or their selves or whatever” (P 10). One lead scientist even evoked a utopian (or Marxian) spirit:

I would like to believe the early socialist and communist writers. That having that extra labour taken care of in the work force will free people up for extra time. That your productivity will increase and therefore the time required for your work will decrease. That has never happened, but I would like to believe it (P 15)

In the course of discussing automation I also asked AI workers whether they thought their own work might be automated. What I learned surprised me and set me on a new path of research.

## 5.13 The Automation of AI Work

My interviewees told me that high-skill work in the AI industry is already being automated – sometimes through the recursive use of AI.<sup>55</sup> While some interviewees expressed concern for their jobs, others told me that they wanted more, not less, automation in their own work. An inference engineer told me that while he thinks his job is safe for now, “[s]omething is going to automate it eventually ... But I’m not worried because I don’t see any kind of like general AI for decades probably” (P 10). Others expressed more imminent threats to their jobs. One data scientist told me:

if I want to secure my job in the future, maybe after five years, I need to learn new technologies. I need to step out of the typical ML models, move toward deep learning or move towards big data ... in my company, a few years before the data cleaning was performed using Excel. But now we have replaced all that with ... complex programs which take in the data and sort out all the anomalies and clean out the data (P 8).

One R&D programmer said:

I think [the chance my work will be automated is] very high ... all work is possible to automate in my opinion ... the last few things to go will be very high end programming, that is, the part of my job that is the most creative and most intellectual ... But most of programming, let’s say 70-80% of programming, is drudgery. It’s a lot of ... moving around code ... doing all sorts of nonsense. All that is going to go away ... the last things that will be to fall are other creative tasks, for instance, poetry and music. But

---

<sup>55</sup> One startup CEO I spoke to was very skeptical of this, though since then the practice has become increasingly prevalent. He said “But ML writ large, things are changing too fast. There’s no point in automating it. You’re going to spend more time doing the automation than you are just doing the work. There’s certain automation that’s very smart. For example, taking snapshots so that you can reuse large chunks of work. Containerization ... There’s a very poor history of code writing code, like code generators. It’s unreadable, it’s not efficient ... It does not work. There are many, many efforts to do this. None in production in real companies. The software that you’re using, the ML stuff that you see when you go to Amazon or whatever. It’s all written by humans. None of it is computer generated” (P 13).



even that, I hesitate to say because there are engines that can prepare ... poetry now ... Maybe mathematics ... I need to start doing math ... Too many people can do [programming]. And machines will soon be able to do it ... I'm actually applying for a math Ph.D. soon (P 4)

One startup CEO discerned a similar threat:

Machines will teach machines better. That's what's going to happen in the future ... I think even ML jobs are going to be really, really replaced ... most software related jobs are going to go eventually. I see them in the next 20 years going ... People will just describe over voice commands – say 'OK, I want to build an Uber-like app that does X, Y, Z'. And then the computer will ... simulate a couple of them ... the idea is not art, what you can build, as funny as it sounds, it's actually finite. It's very large, but it's finite. What's infinite is the design. But in terms of the core functionality what you're really doing is, you're inserting the data, you're applying some logic and you're storing the data. That's finite. What's infinite is, OK I want it to be this colour blue, I want the box here, I want this. And we already have software solutions that do those kind of approximations for us right now ... You know people use Google Firebase so I really see a lot of software jobs going in the future" (P 5).

A similar view was expressed by one data scientist who said "I think the next step for machine learning is to make it possible for anyone to run some pattern recognition or statistic task in little time, get the results for themselves" (P 9). But not all of my interviewees were so concerned. Some expressed quite the opposite opinion. One machine learning scientist told me:

for data scientists and ML scientists in general ... I don't think the fact that it's being automated is a threat. It's sort of a convenience I would say, because then you can go and do more important stuff. More higher level stuff. Things that actually matter ... a few decades ago, you'd sit down and ... use assembly, talking in 0s and 1s to the computer, which was really

hard. And you wouldn't get too far. But now we have very high level languages ... so when Python as a programming language got really popular, you were like, oh! I can do a [inaudible] with so many less characters or lines. But now, even the languages get updated and now the thing that was half a page of code is now just a couple of lines of code. So things go higher ... You can go think about what matters ... So I think it's, at least for the ML community ... not so much of a threat (P 16)

When asked whether he thought his work will be automated, one data scientist said he was looking for a way to automate feature engineering in particular. Another had more general hopes for automation:

I think if you work as a programmer, as a general rule, if you keep doing the same thing for six months you have to start to think of automating. It shouldn't take much longer than that. And if you keep doing the same routine job for a year, you're probably doing it wrong. There must be a higher level of abstraction that can use to go... make things easier for yourself ... you don't have the limitation of tools and material. A civil engineer needs all the steel, and all the workers, to build a thing he designed. But if you design something in the computer you can ... test it, tear it down and build something on top of it that operates it. We should always move towards automating. That's what I always ask everyone in our crew ... what is the routine and if we can automate it, let's do it (P9)

This is a widely held view. A study of 240 data scientists found that “they don't particularly love munging and cleaning data ... [they consider it a] waste of their skills to be polishing the materials they rely on” (Theuwissen 2015, 9). Vorhies (2017b) explains that it “seems like a natural progression to automate what can be safely automated and preserve our time for the creative portions of data science”. McClure (2018) even argues that “automation is the *responsibility* of the seasoned practitioner ... automation ensures our work focuses solely on what is novel to the challenge. When we can rapidly explore the space of possibilities, we promote a much stronger ROI on data-intensive projects”.

However, ease and efficiency are not the only reasons for desiring more automation. A senior software engineer I spoke to explained how the machine learning labour process requires a lot of guesswork and experimentation. This lack of knowledge, he held, was a primary motivation for automating machine learning work:

there's a lot of dark art to the design and layout of a neural network. You're ... drawing a graph and ... you have this expectation that data is going to flow through this graph in some way and you're going to update these things and there's some relationship between the shape of the layout of this graph and the quality of the outputs that come through it. What is that relationship? ... nobody really has more than a cursory understanding of that relationship. Well, we know how to make it really, really bad. We know what not to do ... And if we wanna get better performance we can tweak it a little bit. But we don't actually have a firm model or theory behind it ... if you've got a model going, it's doing these predictions, you're sitting there and you're trying to tweak it to get a little more performance. And some of this tweaking is random and some of it is like enlightened randomness ... And we pay people quite a bit to be better than random at this. But at the end of the day, they're not experts at this. They don't have this deep understanding of how to make the perfect network that works as well as it possibly can ... In my mind, automating this process of exploring, trying different types of networks, that seems a very natural leap to me ... There are many aspects of ML where we actually use ML systems to try to tweak it in various ways. And this is just a growth of that into: can we try to use ML to learn all the parameters of a ML network itself? So it seems like an incremental development to me" (P 12)

In this formulation, the dark art at the heart of machine learning drives the automation of its production. A survey of 240 AI workers agrees that in machine learning there "isn't a codified set of strategies but a wide range of approaches" (Theuwissen 2015, 13). Quoc Le, a Google machine-learning researcher, admits "We do it by intuition" (Simonite

2017b). McClure (2018) agrees that data science is an “art” because “knowing what knobs to turn, and by how much, cannot be codified. It involves experience, conversation with domain experts, and above-all trial and error”. However, it could be argued that this work is being automated *precisely because* it has not been codified. Indeed, I will argue that the automation of machine learning work represents a qualitative shift in automation towards what I call automation without codification. The next section explores this emerging technology in depth.

## 5.14 Automating the Machine Learning Labour Process

My interviewees did not have much detailed information about how exactly machine learning work is being automated, but mention of a technique called hyperparameter optimization put me on the trail of the emerging practice of automatic machine learning or AutoML. AutoML refers to a diverse and growing collection of techniques which automate tasks in all three stages of the machine learning labour process, “from data collection and cleaning, to model development and testing, to production deployment and scaling” (Yao 2019). AutoML has become increasingly prevalent since the time of my interviews. What used to be an experimental project for AI tech giants like Google and Facebook is now the basis for several commodities and open source projects.

In the first stage, data processing, AutoML applications such as MLBox are available for data preprocessing, cleaning and formatting. Zöllner and Huber (2019) survey the variety of applications by which “low quality data can be automatically detected and corrected” although they note that, for now, only basic, general aspects of cleaning are automated and “advanced and domain specific data cleaning is conferred to the user” (19-20). The key task of feature engineering, which requires considerable knowledge about the domain the data represents as well as intuition and a lot of trial and error, is also being automated via “feature learning” networks which can automatically pull out the most relevant features from data sets (Hardesty 2015; Koehrsen 2018).

While so far attempts at automating AI ghost work have not been completely successful, as widespread reliance on human ghost workers evinces, efforts are underway.

Companies such as ClaySciences and DataLoop are producing “automatic annotation” AI

systems which make an initial pass at annotating datasets, so that human ghost workers only have to provide corrects and supplementary information (Singer 2019; Nakashima 2018; DataLoop n.d.). One machine learning expert predicts that in five to ten years the process will be fully automated (Nakashima 2018). The windows provided by the “paradox of automation’s last mile” may not continue to appear forever (Gray and Suri 2019, 17).

AutoML is also being applied in the second stage, where the machine learning model is trained, in at least two ways. The first is model design. Designing machine learning models requires a “significant amount of time and experimentation by those with significant machine learning expertise” because “the search space of all possible models can be combinatorially large — a typical 10-layer network can have  $\sim 10^{10}$  candidate networks” (Zoph and Le 2017). It is often not possible to say what will work best for an unfamiliar problem. Model design is being automated via “neural architecture search” which uses machine learning to design and test thousands of options in the time humans could try a few (Elsken, Metzen and Hutter 2019).<sup>56</sup> Some of these have come up with solutions their creators had previously not considered (Simonite 2017b).

Another AutoML application in the second stage is hyperparameter optimization or “configuring the internal settings that govern the behavior of a model”. It is being automated in a way similar to neural architecture search, via the automatic design and testing of configurations at inhuman speeds (Li and Talkwalker 2018).<sup>57</sup> TransmogrifAI (n.d.) boasts that its AutoML product: “achieves accuracies close to hand-tuned models with almost 100x reduction in time”.

---

<sup>56</sup> Microsoft is employing a similar technique called “human-assisted search” in which “researchers can identify a promising arrangement for massive neural networks, and then the system can cycle through a range of similar possibilities until it settles on this [sic] best one” (Metz 2016a).

<sup>57</sup> Facebook has also developed Asimo, called an “automated machine learning engineer” which “can automatically generate enhanced and improved incarnations of existing models” (Metz 2016b). According to one Facebook scientist: “It cannot yet invent a new AI algorithm ... But who knows, down the road” (Metz 2016b).

Finally, AutoML is also used in the deployment stage. Companies such as Seldon automate the process of “wrapping” a model in a different programming language API while Microsoft Machine Learning Service allows a model to be automatically deployed as a web service, ready for embedding in applications (ThinkGradient 2019).

## 5.15 Automation Without Codification

What is distinctive about AutoML techniques is that they are being applied to aspects of machine learning work which have been resistant to Taylorist codification due to a lack of knowledge which not even expert workers possess. As Vorhies (2017b) writes, the “features being automated require a great deal of skill and experience to get right”. AutoML overcomes or sidesteps the lack of skill and expertise through brute force techniques of experimentation and learning. Automation without codification. Instead of codifying, AutoML enacts a new level of technological recursion. It uses the results of machine learning “to train another machine learning model that can optimize the training of machine learning models” (Metz 2016b).

Classical AI is automation. It eliminates the worker by capturing her knowledge and skills about how to solve a problem (via interviews or time-motion studies). Machine learning can thus be described as “automating automation” (Raschka 2016). Machine learning (ideally) eliminates the need for a worker who first knows the solution to a problem by extracting the solution from data automatically without it first existing in a worker’s head. However, it requires skilled machine learning workers to prepare data, choose algorithms and train models.

AutoML (ideally) eliminates not only the need to know the solution to a problem in advance, but also the need for knowledge of much of the “black art” of machine learning. It replaces aspects of high skill machine learning work with brute force experimentation. AutoML can thus, prolixly, be called “the automation of automating automation” (Mayo 2016). AutoML represents for capital the possibility of dispensing with labour in two senses: in terms of both conceiving of the problem and a solution, as well as understanding and creating the tools by which a solution might be produced

automatically. Ideally, AutoML dispenses with both the need for someone to know how to solve the problem as well as for a skilled creator of an automatic problem solver.

Since I conducted my interviews, AutoML has gone from an experimental technique for the AI giants to play around with to an emergent commodity. It is a primary component of the push for a so-called “democratization” of AI (Steinhoff 2019). This movement is largely motivated by the tech giants for whom it means “mak[ing] it possible for the citizen data scientist (aka lesser or untrained amateurs) to build some models directly” (Vorhies 2017b). But it is also tied up with capitalist competition. Discussing AutoML, Jeff Dean of Google Brain, says: “[c]urrently the way you solve problems is you have expertise and data and computation ... Can we eliminate the need for a lot of machine-learning expertise?” (Simonite 2017a).

It should be noted that AutoML techniques consume huge amounts of computational power. As recently as 2017, experts were saying that AutoML “requires such extreme computing power that it’s not yet practical to think about lightening the load, or partially replacing, machine-learning experts” (Simonite 2017a; see also Simonite 2017b).

However, only two years later in 2019, many AutoML applications are already accessible via cloud platforms. There is now talk of “deep automation” in machine learning (Lorica and Loukides 2018) and the tech giants are working on “end-to-end” or “one-click” AutoML packages which aim to streamline the variegated components of the machine learning labour process into one user-friendly app which will require very little technical knowledge. These include Facebook’s FBLearner Flow, Google Cloud AutoML, Uber Michelangelo ML, Amazon ML, Microsoft Azure and Baidu’s EZDL (Dong 2017). A Facebook engineer describes Flow as a “machine learning assembly line” which will “help engineers build, test, and execute machine learning algorithms on a massive scale ... using as little human grunt work as possible” (Metz 2016b).

Some commentators are skeptical that the whole process can be automated due to a “combinatorial explosion” comprised of interacting “data attributes, models and parameters” (McClure 2018). Others argue that “there is one thing that AutoML definitely cannot do – it never replaces domain expertise” (Yao 2019). Yet, one has to

question the validity of such assertions in light of a recent achievement by Google.<sup>58</sup> The Google AI lab recently demonstrated an end-to-end AutoML system (still in the research phase) which demonstrated “full automation” meaning that “[d]ata and computation resources are the only inputs, while a servable TensorFlow model is the output. The whole process requires no human intervention” (Lu 2019).



**Figure 3: Google AutoML (Lu 2019)**

In April 2019, the Google AutoML participated in a Kaggle Hackathon, competing against grandmaster level data scientists. The task of the competition was “predicting manufacturing defects given information about the material properties and testing results for batches of automotive parts” (Lu 2019). The Google AutoML placed second overall. Google researchers suggest that their success indicates the “potential application of AutoML methods across a wide range of real business problems” (Lu 2019). While it remains unclear how domain expertise might be automated, end-to-end AutoML should not be discounted as hype.

## 5.16 Other Forms of Automation in Machine Learning

Another important note is that AutoML has mainly been applied to supervised learning (Oliveira 2019). How or whether AutoML might be extended to unsupervised learning remains uncertain. But another machine learning technique known as deep learning may present opportunities for the increased automation of unsupervised learning.

Most of the famed successes of deep learning have used supervised or reinforcement learning (Graves and Clancy 2019). However, as one data scientist told me, deep learning “is all about automating feature extraction ... what are the important features that if my

---

<sup>58</sup> One might also cite this statement from Brownlee (2014), proven wrong five years later: “we will probably never have automatic feature engineering”.



system is exposed to it, can best solve the problem? Deep learning basically tackles that part. So it's inherently automation" (P 16). Indeed, the pioneers of deep learning point out that a "key aspect of deep learning" is that the "layers of features are not designed by human engineers: they are learned from data using a general-purpose learning procedure" (LeCun, Bengio and Hinton 2015, 436).

Deep learning (at least theoretically) also addresses the question of domain expertise. According to Runfeldt (2017), "[o]ne of the biggest promises of deep learning (DL) is that domain expertise and handcrafted rules are no longer *required* to create really good predictive or generative models". Instead, DL offers the "possibility of having a single algorithm go from raw data to a desired task, or 'end-to-end learning'" (Runfeldt 2017). While this remains hypothetical as of 2019, a variety of attempts are being made to automate the unsupervised learning, which remains "the Holy Grail of Deep Learning" (Culurciello 2016). One of these is the use of evolutionary algorithms to evolve DL network architectures automatically (Wistuba 2018). Another area of research is meta-learning or "learning to learn" which attempts to overcome the need for massive amounts of context-specific training data by having a system learn its own learning algorithm based on its environment or to automatically "exploit structure in the training domain". Finn (2017) explains: "if we want our agents to be able to acquire many skills and adapt to many environments, we cannot afford to train each skill in each setting from scratch. Instead, we need our agents to learn how to learn new tasks faster by reusing previous experience, rather than considering each new task in isolation".

Another line of research is generative modelling. The generative adversarial network (GAN) has been successfully used to generate very convincing photos of people and objects which never existed in the real world (Goodfellow et al. 2014). GANs do not "simply reproduce the data they are trained ... but rather build a model of the underlying class from which that data was drawn" (Graves and Clancy 2019). Fed enough (unlabeled) images of dogs, a GAN can eventually generate a new image of a dog which it has never seen before. The GAN is comprised of two networks which engage in a game of counterfeiting – one tries to generate new images while the other attempts to discern which ones the first has made. Throughout this process, both networks learn to do their

respective tasks better. Eventually, a GAN can produce synthetic images that are hard to discern from real ones – without human intervention.

A final phenomenon worth noting is the automatic production of synthetic data. One machine learning scientist, who was familiar with adversarial modelling, told me: “I’m aware of a lot of techniques that work on sparse data. Ultimately they do work, don’t get me wrong. But the sort of ideology behind them is always that more data makes them better. It’s very rare that you’ll find anyone in the field who’ll say I’ve got enough data, it’s fine” (P 15). It is not surprising then that the process of automation is also being applied to solve the problem of data availability by creating data instead of functioning with less of it. Research is being conducted on how use machine learning to generate synthetic data instead of gathering it from the real world. One such system, called by its creators the “synthetic data vault” has generated data which data scientists have then trained models on and deployed with accuracy exhibiting “no significant difference” in performance from models trained on real world data (Patki, Wedge and Veeramachaneni 2016, 1).

## 5.17 Conclusion

This chapter has shown how the machine learning labour process is structured by the exigencies of commodity production and is governed by empirical control practices. It has also explored how the stages of the machine learning labour process are being automated through the recursive application of AutoML, which represents a qualitatively novel type of automation – automation without codification. In the next chapter, I analyze *post-operaismo*’s claim for a new autonomy of immaterial labour. I show that if we interrogate the theory of immaterial labour by the light of the labour process analysis of this chapter, claims for a new autonomy seem dubious at best.

## Chapter 6

### 6 The New Autonomy of Immaterial Labour

With the reality of work in the contemporary AI Industry now set out, I am now able to assess *post-operaismo*'s claim for a new autonomy of immaterial labour from capital. First, I clear the ground by noting several substantial criticisms that have been levelled at immaterial labour theory. While these are all serious objections which point to flaws in the conceptual grounding of the theory of immaterial labour, my approach is different. Instead of attempting to reveal theoretical inconsistencies within immaterial labour theory, I test it by comparing it to actual AI work. I do so by reconstructing the technological argument for immaterial labour's purported new autonomy.

While the purported new autonomy of immaterial labour is certainly connected, as Pitts (2018b) argues, to Negri's philosophical shift from dialectics to Spinozan immanence, more important from the point of view of this study is *post-operaismo*'s argument that technological change is at the root of the purported new autonomy of immaterial labour. I show that the three stages of the technological argument for new autonomy are refuted by what I have learned about AI work – immaterial labour *par excellence*. I argue that AI work is, in fact, characterized by its subjection to diverse forms of control by capital and that therefore *post-operaismo* does not provide an adequate theoretical framework for studying AI labour.

Finally, I show how AI work maps better onto the theoretical framework offered by NRM. Because it asserts the continued relevance of Marx's theory of value, NRM allows us to understand the AI Industry as characterized by the same the dynamics of labour and capital which Marx himself noted 200 years ago. The AI Industry today is, in particular, centered around the production of fixed capital, production that itself recursively requires ever larger outlays of fixed capital. A perspective influenced by NRM also allows us to grasp the new type of automation represented by AutoML – automation without codification – as indicative of a new tendency for capital. I suggest that with AutoML capital gains an embryonic autonomy from labour, rather than the other way around. In

conclusion, I argue that what *post-operaismo* thinks is autonomy is in fact obsolescence. However, NRM does not escape this encounter unscathed either. I suggest that NRM's labour theory of value is not prepared to grasp the future trajectory of AI and capital suggested by AutoML.

## 6.1 Extant Criticisms of Immaterial Labour Theory

Before setting out my critique of immaterial labour's purported new autonomy from capital, I will first review some extant criticisms of related elements of the theory of immaterial labour.

A substantial criticism, related to how immaterial labour comprises two incommensurable types of worker, as briefly discussed in Chapter 2, is the claim for the hegemony of immaterial labour. Hardt and Negri hold that immaterial labour has "become hegemonic in qualitative terms and has imposed a tendency on other forms of labor and society itself ... today labor and society have to informationalize, become intelligent, become communicative, become affective" (2005, 109). While Hardt and Negri sometimes hedge this as designating only a "tendency" (2005, 107) they simultaneously hold that there has already been "a real homogenization of laboring processes" (2001, 292). A ready objection is that this hegemony ignores, for instance, the slave labour in African coltan mines and the Fordist manual labour of Foxconn factories. And as Fuchs (2014) asserts, such formulations "hardly account for the continued importance of ... very material resources like oil" (140). Immaterial labour theory suffers from making broad generalizations. As Camfield (2007) argues, positing a "globally hegemonic socio-technical figure of labour in any era in the history of capitalism" both ignores the various forms wage labour has taken over the history of capitalism and glosses over so much detail as to be analytically useless (13-14). This is true, but it does not necessarily preclude the new autonomy of immaterial labour posited by *post-operaismo*.

Two further criticisms have been made of immaterial labour's putative immateriality. The first concerns the necessary materiality of labour. Thompson (2005) points out that labour is "never immaterial. It is not the content of labour but its commodity form that gives

‘weight’ to an object or idea in a market economy” (84). Hardt and Negri (2005) attempt to address such objections by explaining that immaterial labour “remains material – it involves our bodies and brains as all labor does. What is immaterial is its product” (109). However, this does little to fix the issue raised by Thompson – *post-operaismo*’s neglect of the form of value. As Pitts (2017) argues, “[i]f Marx’s theory of value relates not to quantification but to the analysis of form, there is little difference between material and immaterial labours” (333). The same applies to the products of labour. This connects to yet another – most damning – objection to immaterial labour theory.

This is the refutation of *post-operaismo*’s argument against Marx’s labour theory of value. According to *post-operaismo*, the multitude collectively produces, often outside the workplace. Therefore, immaterial labour cannot be quantified, and value can no longer be measured. Marx’s ruminations about the collapse of capital in “The Fragment on Machines” have come true. Pitts (2017) lucidly shows how this position relies on a premise which it simultaneously denies: “postoperaist claims of the Fragment’s realization rest on a disavowed orthodoxy. Despite their professed anti-productivism, they advocate a conventional labour theory of value (LTOV) as a means by which it can be dismissed as historically redundant” (329). In other words, *post-operaismo* posits Marx’s value theory as dead because the distributed immaterial labour of the multitude makes it unmeasurable. However, on a NRM reading of Marx’s value theory, there is already a mechanism built-in to capital for quantifying this distributed sort of labour – the exchange relation, in which value as a form of social mediation, is realized (Pitts 2018a, 11-12). Immateriality is always a property of labour under capital because value is immaterial. Pitts (2018a) therefore argues that *post-operaismo*’s understanding of value is “nowhere near immaterial enough” (6). I will discuss later how value continues to structure the AI Industry and work therein.

*Post-operaismo*’s claim for a new autonomy of immaterial labour has received relatively little attention. An early critique was made by Thoburn (2001) who notes that *Empire* and Negri’s earlier solo work seem “to equate a tendency toward the productivity of communication with an emerging freedom—as if the more fluid and immaterial production becomes, the more it escapes control ... socialized work itself [seems to be]

tending toward autonomy” (86-87). Yet, even as it describes this “emerging autonomy of immaterial labour” it contradictorily ascribes “Foucauldian and Deleuzian conceptions of the immanence of power to all social relations” (Thoburn 2001, 87). Camfield (2007) frames the same theoretical contradiction in terms of subsumption:

Hardt and Negri’s claim amounts to a contention that the real subsumption of labour to capital is retreating, making capital parasitically exploitative of autonomous production. They do not attempt to reconcile this with their contention that the real subsumption of society as a whole to capital has taken place (35)

This critique is substantial. It points out a fundamental contradiction in the theoretical foundations of *post-operaismo*. This, as Pitts (2018b) argues, is a result of Negri’s philosophical shift, which was elaborated by *post-operaismo*, away from Hegel’s dialectics to Spinozan immanence (153). Pitts is correct to point out that the shift to immanence changes the nature of the capital/labour antagonism. However, disputing the philosophical grounds of immaterial labour’s supposed new autonomy is not the goal of this dissertation. Instead, my focus is to assess the adequacy of immaterial labour theory for analyzing the real world. It is logically possible that immaterial labour’s purported new autonomy might obtain whether or not its theoretical foundations are consistent. As I will show, new autonomy is purported to be a result of technological changes in post-Fordism. In the next section, I reconstruct this technological argument.

## 6.2 The Technological Logic of New Autonomy

I argue that *post-operaismo*’s technological logic for a new autonomy of immaterial labour proceeds by three main steps.

Human-machine hybridization > Abstract Cooperation > New Autonomy

The following sections elaborate each stage to reveal the technological logic that underlies ILT.

### 6.3 Human-Machine Hybridization

*Post-operaismo* argues that in post-Fordism the balance of technological power shifts in favour of labour, contrary to the grim depictions of workers becoming the appendages of machines in Chapter 15 of *Capital* Volume 1. Capital increases its organic composition, but “the changing composition of capital ... contributes, subjectively, to strengthening the position of labor” (Hardt and Negri 2017, 114). The “repressive use of technology, including the automation and computerization of production” which capital relied on in previous eras becomes increasingly difficult because technology and labour are no longer at odds with one another (Hardt and Negri 2001, 267). Rather than being subject to increased deskilling and domination, labour gains increasing control over, and merges with, machines: “the hybridization of humans and machines is no longer defined by the linear path it followed throughout the modern period ... [workers] have the capacity to take control of the processes of machinic metamorphosis” (Hardt and Negri 2001, 367). Capital loses control over technology and therefore the logic of increasing organic composition of capital outlined by Marx is supposedly no longer valid. Instead, the multitude becomes increasingly machinic:

The scientific, affective, and linguistic forces of the multitude aggressively transform the conditions of social production ... This consists above all in a complete revision of the production of cooperative subjectivity; it consists in an act ... of merging and hybridizing with the machines that the multitude has reappropriated and reinvented; it consists ... in an exodus that is not only spatial but also mechanical in the sense that the subject is transformed into (and finds the cooperation that constitutes it multiplied in) the machine (Hardt and Negri 2001, 366-367).

Hardt and Negri restate this notion in their latest book: “[w]e should also recognize, perhaps now beyond Marx, as production is increasingly socialized, how fixed capital

tends to be implanted into life itself, creating a machinic humanity” (2017, 114).<sup>59</sup> It is important to note here that while Hardt and Negri certainly draw on the broad and figurative conception of the machine developed by Deleuze and Guattari (see Hardt and Negri 2017, 110) not all of their references to machines can be understood in this register. Deleuze and Guattari’s notion of the machine was deployed to express how human labour is always already enmeshed with technologies, institutions, techniques and various non-humans.<sup>60</sup> In contrast to this transhistorical conception of machine, Hardt and Negri describe a *new* process of human-machine merging, specific to the post-Fordist era. In the past “the productive process ... severely restricted the actualization of the potential that exceeds capital’s bounds ... The affective and intellectual talents, the capacities to generate cooperation and organization networks, the communication skills” (Hardt and Negri 2009, 151-152). But things are different in the era of immaterial labour:

Ever since industrial civilization was born, workers have had a much more intimate and internal knowledge of machines and machine systems than the capitalists and their managers ever could. Today this process of worker appropriation of knowledge can become decisive: it is not simply realized in the production process but is intensified and concretized throughout productive cooperation and spreads throughout the life processes of circulation and socialization (Hardt and Negri 2017, 119).

Technology changes from something alien that confronts the worker in the production process to a collective human prosthesis. Today, “[f]ixed capital ... the memory and storehouse of past physical and intellectual labor, is increasingly embedded in ‘the social individual’” (Hardt and Negri 2017, 114). This new human-machine hybridity process or

---

<sup>59</sup> The incorrect usage of fixed capital here, and in numerous other places, indicates the ongoing lack of interest of *post-operaismo* in the form of value.

<sup>60</sup> “What is critical ... in the concept of the machine, as both Deleuze and Guattari employ it, is a consideration of the organisation of the variety of components in relation to each other that comprises any given machine. For that reason machines can and do exist on any scale and can be both material and immaterial ... visible or invisible ... the key focus was on the assemblage as a collective machine and not ... on the actual technical object” (Savat 2009, 3).



“machinic exodus” defines a new kind of human (Hardt and Negri 2001, 367). Hardt and Negri (2017) explain “when we say fixed capital is reappropriated by laboring subjects we do not mean simply that it becomes their possession but instead that it is integrated into the machinic assemblages, as a constituent of subjectivity” (122).

The human-machine hybridization process is sometimes described as an ongoing struggle. Technologies, the “repositories of social intelligence, become ever more crucial. The curtain raises on the field of battle over control of fixed capital ... Today we must immerse ourselves into the heart of technologies and attempt to make them our own” (Hardt and Negri 2017, 111). While the battle remains to be won, Hardt and Negri (2017) espouse an optimism founded, apparently, on the smartphone use habits of millennials: “When we look at young people today who are absorbed in machinic assemblages, we should recognize that their very existence is resistance” (123). This is a rather perplexing formulation of resistance. I suggest that it becomes comprehensible only if one elaborates an overlooked aspect of *post-operaismo* – its insistence on an irreducible human element in immaterial labour.

*Post-operaismo* evinces a conviction that certain capacities of human labour are necessarily off-limits to implementation in machines. This allows it to affirm the becoming-machine of labour not as capital’s increasing organic composition, but rather as the expansion of labour’s self-directed capacities. Hardt and Negri (2017) assure us that a “Digital Taylorism” can only go so far—while it “sometimes seems as though computer systems, artificial intelligence, and algorithms are making human labor obsolete ... in fact, there are innumerable digital tasks that machines cannot complete” (131). Virno (2003) explains this clearly when he asserts that in post-Fordism, “the means of production are not reducible to machines but consist of linguistic-cognitive competencies inseparable from living labor” (61). In other words, a “decisive role is played by the infinite variety of concepts and logical schemes which cannot ever be set within fixed capital, being inseparable from the reiteration of a plurality of living subjects” (Virno 2003, 106). For Moulier-Boutang, “[w]ithout the power of the living ... which is radically distinct from machinery ... none of this [immaterial labour] can take place” (2012, 163). This power of the living is autonomous creativity (Hardt and Negri

2001, 83) or “collective intelligence, creativity distributed through the entirety of the population” (Moulier-Boutang 2012, 34). For Moulier-Boutang (2012) the “essential point is no longer the expenditure of human labour power, but that of *invention-power* ... the living know-how that cannot be reduced to machines and the opinions shared in common by the greatest number of human beings” (32, emphasis original). This he also defines as “implicit knowledge that is irreducible to machinism, to standardized and codified human capital” (Moulier-Boutang 2012, 54).

## 6.4 Abstract Cooperation

*Post-operaismo* attempts to conceptualize this essentially human creative capacity of the multitude with a reconfigured version of the notion of general intellect presented by Marx in “The Fragment on Machines”.<sup>61</sup> While in Marx’s original formulation, the general intellect refers to capital’s armamentarium of knowledge and capacities excised from human labour and implemented in technology, *post-operaismo* reworks Marx’s notion of general intellect to describe the technologically-enabled mass intellectuality of immaterial labour. In so doing it inverts the concept. Hardt and Negri define general intellect as a “collective, social intelligence created by accumulated knowledges, techniques and know-how” (2001, 364). This definition, in itself, does not necessarily contradict Marx’s formulation. However, when they hold that in post-Fordism “labour is the productive activity of a general intellect and a general body outside measure” (Hardt and Negri 2001, 358) it is evident something has changed. Virno’s (2003) explanation for this change is as follows:

Marx conceives the ‘general intellect’ as a scientific objectified capacity, as a system of machines. Obviously, this aspect of the ‘general intellect’ matters, but it is not everything. We should consider the dimension where the general intellect, instead of being incarnated (or rather, cast in iron) into the system of

---

<sup>61</sup> Reading the notes which comprise *Grundrisse* on par with *Capital* Volume 1 is problematic. As Heinrich (2013) argues, *Grundrisse* lacks the fully developed theory of relative surplus-value which we find in *Capital*. The explosive contradiction that Marx finds waiting within capital in *Grundrisse* is resolved into the difference between absolute value and maximum surplus-value (210-212).

machines, exists as attribute of living labor. The general intellect manifests itself today, above all, as the communication, abstraction, self-reflexion of living subjects. It seems legitimate to maintain that, according to the very logic of economic development, it is necessary that a part of the general intellect not congeal as fixed capital but unfold in communicative interaction, under the guise of epistemic paradigms, dialogical performances, linguistic games (65).

On this account, machines are but one component of the general intellect. For *post-operaismo*, the general intellect refers also, and more importantly, to the new creative capacities of networked labour; it is “one and the same as cooperation, the acting in concert of human labour, and the communicative competence of individuals” (Virno 2003, 65). In other words: “General intellect should not necessarily mean the aggregate of the knowledge acquired by the species, but the faculty of thinking: potential as such, not its countless particular realizations” (Virno 2003, 66). The general intellect of the networked multitude socially produces value for capital through its communicative capacities. Individual workers no longer matter as much as the whole multitude: “today the general intellect is becoming a protagonist of economic and social production” (Hardt and Negri 2017, 114). This new social labour entails a new kind of exploitation:

The dialectic between productive forces and the system of domination no longer has a determinate place. The very qualities of labor power (difference, measure, and determination) can no longer be grasped, and similarly, exploitation can no longer be localized and quantified. In effect, the object of exploitation and domination tend not to be specific productive activities but the universal capacity to produce ... abstract social activity and its comprehensive power (Hardt and Negri 2001, 209).

*Post-operaismo* thus posits a broadening and softening of exploitation. Broadening because all social existence is now productive labour and therefore exploitation is a net cast wide across society. Softening because there is no definite moment of exploitation – it is “omnipresent” but “there is nothing miserablism about it” (Moulier-Boutang 2012,

92). Exploitation is not eliminated, just redefined as “the expropriation of cooperation and the nullification of the meanings of linguistic production” (Hardt and Negri 2001, 385) or the “capturing positive externalities” (Moulier-Boutang 2012, 55). Capital can only appropriate the results of social production, not direct its activities. Capital “is increasingly external and has an ever less functional role in the productive process” (Hardt and Negri 2009, 142). Capital can therefore only survive by becoming parasitic on the products of autonomous production: “capital expropriates cooperation ... at the level of social production and social practice” (Hardt and Negri 2009, 140-141). Exploitation becomes “the private appropriation of part or all of the value that has been produced as common” (Hardt and Negri 2005, 145).

On this basis, Hardt and Negri (2017) discern “a challenge or even a potential threat to capital because the primary role in the social organization of production tends to be played by the living knowledges embodied in and mobilized by labor rather than the dead knowledges deployed by management and management science” (115). Immaterial labour exhibits a novel form of cooperation which Hardt and Negri (2001) call “abstract cooperation” (296). This is “*completely immanent to the laboring activity itself*” while in previous eras of capitalism, cooperation had to be organized by the capitalist controlling the labour process (Hardt and Negri 2001, 294, emphasis original). While capital implemented scalable cooperation when it imposed the division of labour, labour is now “immediately a social force animated by the powers of knowledge, affect, science, and language” (Hardt and Negri 2001, 358). This is enabled by “computerization” (Hardt and Negri 2001, 292). Networks generate a novel way of working or a “becoming common” (Hardt and Negri 2005, 129). The “assembly line has been replaced by *the network* as the organizational model of production” (Hardt and Negri 2001, 295, emphasis original). For Lazzarato (1996), “immaterial labor constitutes itself in forms that are immediately collective, and we might say that it exists only in the form of networks and flows” (154). The networked multitude engages in abstract cooperation, employing communicative and creative capacities which capital cannot implement in machines. Capital’s drive towards an increasing machinic state thus runs into an obstacle and labour gains a new autonomy from it.

## 6.5 New Autonomy from Capital

*Post-operaismo* thus posits an emergent capacity for “autonomous production” for immaterial labour (Hardt and Negri 2001, 276). This is also formulated as “real (and increasing) productive powers and capacities for autonomy” (Hardt and Negri 2017, 77) as well as the “radical autonomy of the productive synergies of immaterial labor” (Lazzarato 1996, 140). This new autonomy is described as a result of machinic hybridization and abstract cooperation in a wealth of grandiose proclamations: “[t]o the same degree that capital, as this process [of hybridization] proceeds, loses the capacity for self-realization, the social individual gains autonomy” (Hardt and Negri 2017, 114). Workers, “having incorporated the tools of production, having been metamorphosed anthropologically, act and produce machinically, separately and autonomously from capital” (Hardt and Negri 2017, 133). Immaterial labour thus “becomes in cooperation increasingly abstract from capital—that is, it has a greater ability to organize production itself, autonomously, particularly in relation to machines” (Hardt and Negri 2017, 117).

The “productive social cooperation of workers endowed with fixed capital ... poses the potential for the autonomy of workers, inverting the relation for force between labor and capital” (Hardt and Negri 2017, 115). Labour has “reached such a level of dignity and power that it can potentially refuse the form of valorization that is imposed on it and thus, even under command, develop its own autonomy” (Hardt and Negri 2017, 117). Capital becomes increasingly machinic, but it cannot use machines to increase relative surplus value extraction any longer because production now depends on processes of autonomous communicative cooperation by the multitude. In this development, *post-operaismo* discerns the kernel of a future beyond capital.

Autonomous production “exceeds capitalist relations ... grants labor increasing autonomy and provides the tools or weapons that could be wielded in a project of liberation” (Hardt and Negri 2009, 137). This is not only a liberation from work.

Autonomy from capital entails a revolutionary transformation of all life. Lazzarato sees “a ‘silent revolution’ taking place within the anthropological realities of work and within the reconfiguration of its meanings” (Lazzarato 1996, 140). But Hardt and Negri (2001) suggest that immaterial labour’s new autonomy “seems to provide the potential for a kind of spontaneous and elementary communism” (294). Workers, “[h]aving incorporated the productive tools and knowledges into their own minds and bodies ... are transformed and have the potential to become increasingly foreign to and autonomous from capital. This process injects class struggle into productive life itself” (Hardt and Negri 2017, 115). Emerging is a “new productive nature ... a new form of life that is at the base of a new mode of production” (Hardt and Negri 2017, 119). The sense of autonomy that comes with this new mode of production is qualitatively novel in the history of capitalism:

Is this autonomy the same as the forms of worker autonomy we spoke of in earlier phases of capitalist production? Certainly not, because now there is a degree of autonomy not only in regard to the processes of production, but also in an ontological sense—labor gains an ontological consistency, even when still completely subordinated to capitalist command (Hardt and Negri 2017, 117)

Because production now takes place beyond the workplace, labour “appears simply as the power to act” (Hardt and Negri 2001, 358) or as “collective human activity as world-constituting practice” (Pitts 2018b, 160). This new ontological autonomy is thus equivalent to an immanent “potentiality of the multitude” to create (Hardt and Negri 2001, 82; Pitts 2018b, 153-160). Capital cannot emulate this potentiality and thus can only survive on what scraps it can steal from the autonomous production of immaterial labour.

The technological argument for a new autonomy of labour can be summed up in one quote from Hardt and Negri (2009): “the powers of the new technical composition of labor-power cannot be contained by the capitalist modes of control” (143). The proliferation of ICTs and technical skills in workers enables a machinic hybridization of immaterial labour (which is inherently resistant to total machinic takeover by capital).

Hybridization produces a new social subject (multitude) which possesses a new capacity for self-organization or abstract cooperation. This technologically-enabled cooperation cannot be implemented in machines and so capital can only parasitically capture value produced by the multitude. Immaterial labour thus produces increasingly autonomously from capital.

Before assessing the theory of immaterial labour, it is important to mention the ambiguity with which claims for it are formulated. *Post-operaismo* tends to be unclear as to whether the revolutionary changes it describes have already occurred, are underway, or might occur in the future. As the quotations above show, immaterial labour is often described as, for instance, fully hegemonic, in other passages we are told that some “propositions have to be understood as indicative of a tendency” (Hardt and Negri 2017, 118).

Likewise, while labour is sometimes depicted as beyond the control of capital, at other times it is described as possessing only “certain limited margins of autonomy” (Hardt and Negri 2017, 133) or capital’s influence over labour is described as holding, but “not to the same extent” (Hardt and Negri 2009, 140). Sometimes capital is described as solely parasitic on autonomous production, while in other formulations, labour remains “subordinated to the mechanisms of the extraction of value by capital” although they no longer function properly (Hardt and Negri 2017, 117).

In the interest of producing the most productive engagement possible with immaterial labour theory, I want to represent the argument for a new autonomy of immaterial labour as strongly as possible. This is known as creating a “steel man,” in contrast to a “straw man,” which is an easy-to-defeat misrepresentation of an argument (lukeprog 2011). It is certain today that labour has not achieved full autonomy from capital, so I will abandon that interpretation immediately. I suggest the following reconstruction of *post-operaismo*’s technological argument for new autonomy as a cyclical process comprising three stages. This formulation has the benefit of reining in *post-operaismo*’s more far-fetched statements.

#### The Technological Argument for the New Autonomy of Immaterial Labour

Machine hybridization > Abstract Cooperation > New Autonomy

Summarily: machinic hybridization enables abstract cooperation which enables a new degree of autonomy for immaterial labour. Whatever level of autonomy has been obtained in the course of this process is then used to appropriate more technology, augment hybridity and expand the capacities of abstract cooperation, beginning the process over again. Eventually, labour achieves full autonomy and escapes from capital, which then perishes, starved of surplus-value. On this reading, immaterial labour's new autonomy is currently partially achieved and is set to increase in the future.

## 6.6 AI Work and Immaterial Labour Theory

The study of AI work conducted in this dissertation does not support the technological argument outlined above. AI work is far from exhibiting the novel human-machine hybridization, abstract cooperation or autonomy of labour posited by immaterial labour theory. On the contrary, work in the AI Industry continues to be structured by the exigencies of capital valorization; it evinces capital's "immanent drive, and ... constant tendency, towards increasing the productivity of labour, in order to cheapen commodities and, by cheapening commodities, to cheapen the worker himself" (Marx 1990, 436-437). AI work exhibits characteristics of labour under capital which date back to Marx's time.

The theory of immaterial labour is founded upon the becoming-machine of labour as the multitude. According to this proposition, immaterial labour wrests control over machinery from capital, begins to use it for its own ends and even merges with it. Capital's affinity with machines is reversed. Yet, it seems to be the case that AI work relies essentially on machines developed and controlled by capital. While AI research continues in academia, many important breakthroughs and all applied AI technologies are produced by capital. And, as I have shown, AI research in general and the AI Industry in particular, would not be possible without the huge amounts of computing power provided by the cloud. This necessary infrastructure for AI work is controlled by a handful of ruthlessly competing AI tech giants. In Chapter 4, I discussed how the originally non-profit think tank OpenAI recently converted to a for-profit model, supported by funding from Microsoft, who is now their sole cloud provider. OpenAI justifies this change by



simply stating that AI work “requires a lot of capital for computational power” (Brockman 2019). If an AI think tank funded by the philanthropy of billionaires like Elon Musk cannot operate outside of capital, what hope is there for a proletarian AI initiative?

One might, however, argue that considerations of hardware are misdirected and that the proliferation of open source AI software tools signifies a more important phenomenon of machinic hybridization. Both the AI Industry and *post-operaismo* are united in their enthusiasm for open source software (Moulier-Boutang 2012, 79-83). I have shown in Chapter 4, however, that companies in the AI Industry are using the open-source model merely as a new weapon in ongoing attempts at market domination. Indeed, the notion that open-source software is developed by independent users might be wholly overturned as the AI Industry takes off. While equivalent statistics for the AI Industry are not available, the 2017 Linux kernel received “well over 85 percent” of its contributions from “developers being paid for their work” (Corbet and Kroah-Hartman 2018, 15). With the AI tech giants benefitting from open source development, the practice may soon be just another capitalist labour cost-saving technique.

There is a third sense in which *post-operaismo*'s notion of human-machine hybridity is contested by AI work. Contrary to *post-operaismo*'s assertions of a dramatic reversal in control over technology, my analysis of AI work shows the continuation of the same processes of deskilling, fragmentation and automation that Marx and Braverman noted in earlier types of labour, as Chapter 5 details. As an NRM perspective, for which the determination of the value form continues to be efficacious, would expect, all levels of AI work are already being turned over piecemeal to machines, indicating the increasing capacities of capital's general intellect.<sup>62</sup> From the high-skill work of data scientists and machine learning engineers to the digital manual labour of the ghost workers who annotate training data, AI work is being degraded in the Bravermanian sense and the

---

<sup>62</sup> Terranova (2004), who endorses a version of immaterial labour theory, argues that *post-operaismo*'s humanist reconfiguration of the general intellect is justified because, if it were not “the Marxian monster of metal and flesh would just be updated to that of a world-spanning network, where computers use human beings as a way to allow the system of machinery (and therefore capitalist production) to function” (87). This, I suggest, is precisely the trajectory indicated by my analysis of the AI Industry.

organic composition of the AI Industry continues to increase. Gray and Suri (2019) argue that “the desire to *eliminate* human labor always *generates* new tasks for humans” (17, emphasis original). However, AutoML, with its novel capacity for automation without codification, and in particular end-to-end AutoML, with its aim to eject the human element entirely from the production of machine learning models, may represent a new type of machine-labour dynamic which meets the needs of capital in an entirely new way. I return to this point below.

Overall, *post-operaismo*'s notion of human-machine hybridity is founded upon the belief that certain social and creative capacities of humans simply cannot be achieved with machines, and that therefore human labour is safe from widescale replacement in production. Here *post-operaismo* agrees with capitalist centaur theorists. But neither centaur theorists nor *post-operaismo* supply any support for this claim. A brief look at the history of technology shows that assuming an eternal divide between what machines can and cannot do has yet to work out. To take one example, Dreyfus (1972) argued for inherent limits to AI which have since been overcome by the ascendant machine learning approach. Today, AI research continues to advance. The “impossible-to-automate human capacities” demarcated by *post-operaismo* are already being automated (Dyer-Witheford, Kjøsén and Steinhoff 2019, 65). I therefore argue that quite the opposite to the multitude's becoming-machine seems to be happening. Instead, capital, in the form of machinery, is taking on increasingly human capacities, even if it is unlikely to obtain subjectivity any time soon. In other words, AI work evinces the “development of the means of labour into machinery ... the historical reshaping of the traditional, inherited means of labour into a form adequate to capital” (Marx 1993, 694).

Immaterial labour theory also relies on a central claim for abstract cooperation, by which the multitude produces socially, outside of capitalist control. I have shown, on the contrary, that AI work is characterized by a combination of decentralized control practices and software, such as the Scrum methodology and the JIRA software, which are applied by management. Cooperation is still instituted by capital, in accord with its own ends. While AI workers have a limited sense of what Barret (2005) calls “technical autonomy” their work is surveilled and optimized for competitive advantage through a

variety of social and technical mechanisms (82). What Meiksins (1996) says about franchise models applies also to development methodologies like Scrum – though ostensibly decentralized, they ultimately serve to “eliminate local autonomy and to maximize the degree of control from the center” (156). Cooperation in AI work remains directed by the exigencies of commodity production. As Pitts (2018b) puts it, “[c]apitalism is characterized by categories of social mediation. They persist regardless of whether a work uses a keyboard or a hammer, ideas or nuts and bolts. And in this is implied the persistence of means of measure and time discipline familiar to the pre-‘social’ factory” (loc 4002). Software development methodologies are merely a new way of measuring and disciplining the productivity of labour.

*Post-operaismo* posits a new autonomy of immaterial labour but AI work evinces nothing of the sort. This is fundamentally because, AI work, like all work under capital, continues to be structured by the value form and is thus characterized by precisely the same dynamics which Marx noted in studying industrial labour two hundred years ago, even if different tools and management techniques are used. Capital’s inherent drive towards an increasingly machinic state continues in AI work. I must agree with Camfield (2007) who asserts that the notion of immaterial labour taking place outside of the circuits of capital is “little more than an example of wishful thinking” (47). AI work instead lies at the heart of cybernetic capital. Not only is AI a type of fixed capital that is finding more uses daily. As I, and my co-authors, have argued elsewhere, capital is positioning AI as a new intelligent infrastructure, or in Marxist terms, as a new part of the “general conditions of production” (Steinhoff 2019; Dyer-Witheford, Kjøsén and Steinhoff 2019, 30). The producers of this new infrastructure, the AI Industry, are, like any other capitalist industry, subject to the inexorable drive towards greater automation. I thus disagree also with the LPT thinkers who have described software work as a “craft” which resists deskilling and automation because to “fully standardize computer programming ... would require the seemingly omniscient knowledge of both the emergent problems and the associated solutions” (Andrews, Lair and Landry 2005, 67). I have shown how with machine learning the need for knowledge of a solution is, at least ideally, dissolved. And with AutoML, an even further abstraction from the problem is possible. Hardt and Negri (2017) argue:

Today we can really begin to think of a reappropriation of fixed capital by the workers and the integration of intelligent machines under autonomous social control into their lives, a process, for example, of the construction of algorithms disposed to the self-valorization of cooperative social production and reproduction in all of their articulations (119)

However, the production of intelligent machines remains under the control of capital. The production of AI has been automated, and that process of automation is itself being increasingly automated with the recursive technique of AutoML, the significance of which Caffentzis (2013) adumbrates when he speaks of the “automatization of automation” (129). I suggest that AutoML, with its capacity for automation without codification, signifies, in embryonic form, a trajectory for capital that few extant Marxist theories of labour, capital and machines are prepared to grasp: the “complete dissociation of living labour ... from the production process” (Ramtin 1991, 58).

## 6.7 Autonomy of What?

*Post-operaismo* holds that because of the new autonomy of immaterial labour, capital will resort to desperate strategies which will generate new and fiercer social antagonisms than those of previous eras. Hardt and Negri (2017) assert that if “capital can expropriate value only from the cooperation of subjectivities but they resist that exploitation, then capital must raise the level of command and attempt increasingly arbitrary and violent operations of the extraction of value from the common” (123). Lazzarato (1996) agrees that immaterial labour “re-poses the [class] antagonism at a higher level, because it both mobilizes and clashes with the personality of the individual worker” (135). This position, however, assumes the ongoing irreducibility of a human component of labour, which as I showed above, is uncertain, if not dubious. One might argue, however, that the recent spate of activism in the tech industry is a sign of some sort of autonomy of labour emerging in this sector of capital. I will take up this possibility in the conclusion.

But first, I want to advance the argument that the automation without codification enabled by AutoML is indicative of a possible tendency for cybernetic capital – a technique for augmenting capital’s general intellect *without* first capturing knowledge or

skills from human labour. This, I suggest, casts today's AI Industry as the birthplace not of newly autonomous immaterial labour, but, conversely, of a technology by which capital could become increasingly autonomous from human labour. While the notion perhaps seems excessively futuristic, as I mentioned in Chapter 2, the possibility was considered as early as the 19<sup>th</sup> century by David Ricardo, who reasoned that “[i]f machinery could do all the work that labour now does, there would be no demand for labour. Nobody would be entitled to consume anything who was not a capitalist, and who could not buy or hire a machine” (Ricardo 1951-1973, VIII: 399-400, cited in Kurz 2010, 1195). Here Ricardo imagines a capitalism in which all workers have been replaced with machines while human capitalists continue to direct their operations and reap the spoils of machine-labour. But there is also the possibility that capital could become “autonomous not just from human labour, but from human beings tout court” (Dyer-Witford, Kjølseth and Steinhoff 2019, 139).

This scenario was first fleshed out by Nick Land.<sup>63</sup> Starting in the 1990s, Land championed a theory of “capital autonomization” (Land 2018). Essential to capital autonomization is his notion of a “teleological identity of capitalism and artificial intelligence” (Land 2014). He theorizes capital as a process of positive (i.e. recursive) cybernetic feedback which employs computers, and then AI, to amplify and accelerate its feedback loops (Land 2012, 286-300). For Land, capital is software which has historically run on human hardware. The market was an early form of its intelligence. But with the proliferation of computers, Land suspected that capital was in the process of switching to a new, more adequate substratum – machines. Land (2012) asserts that capital “only retains anthropological characteristics as a symptom of underdevelopment ... Man is something for it to overcome” (446). He theorizes AI as capital's method of escape from its reliance on human labour: “[j]ust as the capitalist urbanization of labour abstracted it in a parallel escalation with technical machines, so will intelligence be transplanted into the purring data zones of new software worlds in order to be abstracted

---

<sup>63</sup> I wish to emphasize my disdain for Land's more recent neo-reactionary, alt-right output. I draw here only on his early theorization of cybernetic capital.

from an increasingly obsolescent anthropoid particularity” (Land 2012, 293). For Land (2012), capital tends towards a purely cyberspatial form where it can function, without human intervention, at supercomputing speeds; capital institutes a “becoming inhuman of cognition, a migration of cognition out into the emerging planetary technosentience reservoir ... where human culture will be dissolved” (293). More recently, Krašovec (2018) develops a similar theory of cybernetic capital, asserting that “anthropocentric theories of capital” are unable to grasp how “autonomous machines and artificial intelligence” are moving “towards an ever greater independence of capital from humanity” via an “autonomisation that simultaneously denotes self-referentiality”. He suspects that we are entering a time in which:

human labour and intellect are becoming, from the point of view of capital, increasingly cumbersome, inert and obsolete and thereby redundant, a time where technologies of design, production and multiplication of technological innovation are immanent to capital itself (and are not borrowed from humanity) (Krašovec 2018).

If this were to occur, capital would become a self-augmenting automaton whose sole concern (value) would be expressed in competition-driven continual upgrades of its production apparatus: “production results in profit, which provides the possibility for improvement, a technological upgrade of the process of production and so on into infinity” (Krašovec 2018). This, he poses in counterpoint to anthropocentric theories of capital, as the notion of the “real autonomy (RA)” of capital, or a:

technological dynamic that is regulated and determined by competition. In the phrase RA we have ‘autonomy’ because this logic is non-human, it is independent of human intentions and/or needs, and ‘real’ because this is actual autonomy, not a fetishistic illusion, it is not attribution of mystical intrinsic characteristics to things (to money or machines, for instance), but a description of how capital relation actually functions (Krašovec 2018)

There are problems with Krašovec’s analysis (his understanding of the value form), but he usefully supplies a term for the tendency indicated by AutoML’s automation without

codification – the real autonomy of capital, not labour.<sup>64</sup> Even if the real autonomy of capital is a concern only for future generations, the tendency I have noticed in AutoML is significant for Marxist theory today – and not only for *post-operaismo*. As I indicated in Chapter 1, the long-term possibility of increasingly autonomous capital posed by AI also entails a problem for NRM.

In the face of the possibility of the real autonomy of capital, any theorization of capital and labour must question the long-term viability of the Marxist dictum that “[t]here are no powers of capital that are not ultimately the collective social powers of labour (or the powers of nature, machinery and science mobilized by collective social labour)” (Smith 2009, 124). It is undoubtedly important for Marxist theory not to attribute capacities to machines which they do not possess. Hardt and Negri (2017) are correct to point out how algorithms can often obscure the “potent figure of labour” that goes into producing them (118). One must ask at what point is a machine produced by machines no longer a product of human labour? How many generations of AI, autonomously produced by AI, autonomously produced by different AI, etc., will it take before the distant origin in human labour becomes negligible? This scenario, which I suggest might be called the “cybernatural,” is beyond the theoretical reach of any Marxist theory which cannot imagine the possibility of machines becoming exploitable in the technical Marxist sense. A value form perspective such as NRM, however, has some advantages here because it emphasizes the already-inhuman nature of capital, which thinks only in terms of value.<sup>65</sup> Marx (1992) writes that it is “only the function of a product as means of labour in the production process that makes it fixed capital. It is in no way fixed capital itself” (240).

---

<sup>64</sup> This example fails to distinguish between the forms of fixed and variable capital and thus attributes labour to a simple app which has certainly not been proletarianized: “a hired programmer writes a code for an application that offers yoga advice, let’s say. A few extra people handle the marketing and promotion of the application, but the app does most of the work by itself: it answers the questions of consumers, adapts to situations, recalls previous queries etc. And in the end, the company earns profit, so the activity must have been productive and brought surplus value, which means that we have a situation where in capitalist economic activity it is actually the (flexible and intelligent) app that is being exploited” (Krašovec 2018). The proletarianization of machines is not a topic that can be adequately addressed here. For an elaboration on this topic, see Dyer-Witthford, Kjösen and Steinhoff (2019).

<sup>65</sup> On this point, I am indebted to Kjösen (2013), who first elaborated it.

Marx explains that any given object under capital has its natural form, as well as its social form – which is to say, value (Marx 1990, 138). A human, as Dyer-Witthford, Kjøsen and Steinhoff (2019) explain, can therefore be made into fixed capital – indeed, this is the economic form of slavery: “The slave ... has the same ontological status – appearing in the form of fixed capital – as machines or animals when used in a capitalist production process” (135). It is possible that the inverse of this ontological change could occur with sufficiently advanced machines. AI functioning as fixed capital could become variable capital, or in other words, be proletarianized and exploited, and enable the harvest of surplus-value (Dyer-Witthford, Kjøsen and Steinhoff 2019, 137-138). While NRM is better prepared to grasp this scenario insofar as its emphasis on the form of value lacks a notion of irreducible humanity, few NRM thinkers would be ready to endorse this notion since it demands a fundamental revision to the concept of labour; labour would no longer be an exclusively human capacity. Thus, even as this dissertation uses NRM to critique *post-operaismo*, it recognizes that NRM is also not fully prepared to theorize the dynamics of labour and machine under advanced cybernetic capitalism. While this dissertation cannot offer an adequate formulation, it indicates the need for Marxist theory of the future to not underestimate the inhuman nature of capital’s nihilistic drive towards the machinic.

Not only dystopian sci-fi inflected analyses come to dire conclusions when assessing the increasing organic composition of cybernetic capitalism. Other more immediate analyses not enamoured of immaterial labour pronounce analogously grim expectations for the future. The collective Endnotes comes to similar conclusions based on their reassessment of Marx’s discussion of surplus populations. Marx (1990) suggests that a result of capital’s machinic tendency is the production of a “redundant working population... which is superfluous to capital’s average requirements for its own valorization” (782). This “surplus population” is described as a cyclical by-product of industrial production, or a sort of pressure-release valve for capital (Marx 1990, 517). *Endnotes* has, however, argued that the long-term dynamics of capital’s becoming-machinic compel a secular increase in surplus populations, not merely cyclical crises. They hold that “*labour-saving technologies tend to generalise, both within and across [production] lines, leading to a relative decline in the demand for labour*” (Endnotes 2010, emphasis original). And while



this relative decline has been counteracted throughout the 20<sup>th</sup> century by various countervailing forces, the ultimate direction must be towards an absolute decrease in necessary labour:

Marx ... notes that the higher the organic composition of capital, the more rapidly must accumulation proceed to maintain employment, “but this more rapid progress itself becomes the source of new technical changes which further reduce the relative demand for labour.” This is more than just a feature of specific highly concentrated industries. As accumulation proceeds, a growing “superabundance” of goods lowers the rate of profit and heightens competition across lines, compelling all capitalists to “economise on labour”. Productivity gains are thus “concentrated under this great pressure; they are incorporated in technical changes which revolutionise the composition of capital in all branches surrounding the great spheres of production (Endnotes 2010).

Due to its recursive technological revolutionizing, capital “produces a relatively redundant population ... which then tends to become a consolidated surplus population, absolutely redundant to the needs of capital” (Endnotes 2010). Today, as swaths of desperate migrants throw themselves against hostile borders and venture across seas in leaking rafts, as conventional labour diminishes while platform labour and ghost work proliferate and as new applications of automation continue to appear, it is difficult to deny this analysis plausibility.

While Endnotes is not predicting a sudden jobs apocalypse, but rather describing an on-going dynamic of capital which meets with countervailing tendencies, it is possible to read Marx’s musings on surplus populations as a sort of alternate, dark ending to the high-tech scenario he explores in “The Fragment”. Pushing the logic of the surplus populations argument to the limit, one can imagine a scenario in which a hyper-automated capital, which has eliminated its reliance on human labour, does not implode due to a vacuum of value, but instead, with proletarianized machines allowing the continued harvest of surplus-value, jettisons its former human substratum completely,

creating a global population superfluous to its needs. Here there is no capital parasitic on the autonomous production of immaterial labour. On the contrary, here is an autonomous capital with which destitute humans, for whom a wage is impossible, must parasitically engage to obtain the necessities of survival.

Of course, it remains uncertain how widely automation without codification might be applied. But if it does find wide applicability, as machine learning increasingly is, we can expect a radical new wave of real subsumption. This is one way in which the “ground up disruptions” which reduce whole industries to a few workers, as mentioned in Chapter 4, could occur (K.F. Lee 2018, 177-178). At the very least, the notion of automation without codification should remind us that capital is not “exclusively a reorganisation of human production” it can also be a “radically new, alien way of production” (Krašovec 2018). In the striking words of Smith (2009), capital is a “higher-order alien power operating at the level of society as a whole. It systematically selects for human ends compatible with its end, ‘the self-valorization of value’, and systematically represses all human ends that are not compatible with this non-human end” (123). AI, machine learning and AutoML function as new prostheses for capital’s reconfiguration of society.

## 6.8 Conclusion

Immaterial labour theory has not fared well in its assessment by the example of the AI Industry. I have shown that AI work, which represents a prime example of immaterial labour, does not exhibit the new autonomy from capital attributed to it by *post-operaismo*. AutoML was the culminating moment in an exploration of AI and capital and their conjunction in the contemporary AI Industry. Both AI (and computing generally) as well as capital, exhibit properties of recursion which make suggestions that they possess a “teleological identity” plausible (Land 2014). This investigation thus concludes that what *post-operaismo* theorizes as a new autonomy of labour is quite the opposite – the growing obsolescence of labour for capital and the growing autonomy of capital from labour.

Capital, as Marx recognized, is always searching for ways to overcome barriers to its valorization. The circulation of money, for instance, allows capital to overcome the temporal and spatial barriers inherent in direct exchange (Marx 1990, 209). Capitalist

machinery enables the overcoming of a different barrier. The “automatic mechanism” helps “reduce to a minimum the resistance offered by man, that obstinate yet elastic natural barrier” (Marx 1990, 527). AI-powered cybernetic capitalism requires a theory of labour, capital and machine that can grapple with increasingly autonomous kinds of capitalist machinery which encroach increasingly on capacities historically reserved for humans. *Post-operaismo* does not offer such a theory.

## Chapter 7

### 7 Conclusion

This dissertation has assessed *post-operaismo*'s claim for a new autonomy of immaterial labour from capital through an analysis focused primarily on high-skill data science work in the AI Industry. I have shown how the history of the AI Industry and the form it takes today are characterized by an increasing technological recursion which allows capital to substitute a mounting array of machines for the capacities of human labour. I have argued that AI work is characterized by capitalist control in ways familiar from previous industries. The analysis of the labour process presented here shows that both Taylorist methods of control applied since industrial capitalism and more recent forms of “empirical” control via Agile and Scrum development methodologies are present in AI work (Schwaber and Beedle 2002, 25). Further, I have explored how, with the emergent techniques of AutoML, AI work is itself being automated through the recursive application of AI. I have argued that, antipodally to the purported new autonomy of immaterial labour claimed by *post-operaismo* theorists, the possibility of automation without codification presented by AutoML indicates a technological means by which capital might increase its autonomy from labour.

With the argument complete, this concluding chapter discusses some limitations of this dissertation, ways they might be addressed in future work, and some considerations on how critical thought and practice concerning AI in general might move forward. But first, I want to raise one critical point which has been implicit throughout this dissertation, although not explicitly argued. This is the importance of considering AI as a technology of automation.

#### 7.1 The Importance of Automation

There are many possible ways to conceptualize AI, as Chapter 1 shows, and not all of them consider automation an important aspect. In his illuminating Foucauldian archaeology of machine learning, Adrian Mackenzie (2017) recognizes the technology as a form of control, but is “less certain about treating machine learning as automation”

because “[l]earning from data ... often sidesteps and substitutes for existing ways of acting, and practices of control, and it thereby reconfigures human-machine differences” (8). Rather than automating, Mackenzie (2017) argues, machine learning reconfigures various objects and populations in an operation he calls “probabilization” (105). This means that machine learning “is not simply automating existing economic relations or even data practices” (Mackenzie 2017, 13). He thus argues that to

qualify or specify how machine learners exist in their generality, we would need to specify their operations at a level of abstraction that neither attributes a mathematical or algorithmic ideality to them nor frames them as yet another means of production of relative surplus value (Mackenzie 2017, 17)

This dissertation, however, demonstrates that any critical analysis of AI is inadequate if it does not consider how AI is produced today as a commodity by powerful capitalist firms and is increasingly deployed as fixed capital in diverse sectors of capital. Here I have to agree with AI luminary and entrepreneur Andrew Ng (2017) who suggests that if you are trying to understand AI’s near-term impact, “don’t think ‘sentience.’ Instead think ‘automation on steroids’”. Indeed, Mackenzie’s book presents a wealth of critical insight, but does not suggest that it might be relevant that machine learning sits at the center of a burgeoning industry. This is not to say Mackenzie is wrong, only that he is missing part of the picture. I have suggested that machine learning might be *both* automation and a way of sidestepping or substituting for existing ways of control. According to my argument, AutoML precisely represents the nascent possibility of automating *by* sidestepping the capture of human skills and knowledge. If this is so, then more rather than less attention to conceptualizing AI as automation ought to be paid because the meaning of the term may be changing at a fundamental level.

## 7.2 Limitations and Ways to Address Them

This dissertation has many limitations. Some of these concern the analysis of work in the AI Industry while others pertain to the theoretical argument developed on the basis of this analysis.

One particularly significant limitation concerning the analysis of work in the AI Industry is the relatively small sample size of interviews. While this was necessary due to the limited time and resources available, a larger sample would ideally be able to capture more diverse perspectives and allow for a comparative analysis between the different types of AI workplace, from tech giant to startup. One particularly notable lacuna in the interviews is the lack of female interviewees. Another is the lack of ghost workers. As I have discussed in Chapter 4, the AI Industry and larger tech industry are dominated by white males. The experience of female AI workers would be an asset in developing a better critical perspective on AI work as well as arguments about new hegemonic forms of labour such as those advanced by *post-operaismo*.

The study could also have benefitted from interviews with the purchasers of AI commodities. This dissertation focused on the production of AI. However, as NRM theorists have emphasized, production does not function as valorization unless the commodity is sold, and its value is realized. Therefore, a complete analysis of AI production would need to include the distribution of AI commodities.

Another significant limitation of this study is its lack of workplace observation of the labour process. The understanding of the labour process presented could be greatly enhanced by ethnographic observation of AI workers as they go about their workday. This research would draw on work by Forsythe (2001) who pioneered the anthropological study of AI researchers. There are undoubtedly social interactions, attitudes and other behaviour which were not addressed by the interviews but could contribute to better understanding what goes into the production of AI products.

Further, it could be beneficial to study AI products themselves. How do AI workers, as Brian Brown (2012) asks of Flickr contributors, impart their subjectivities to the products of their labour (126-127)? This question becomes especially interesting with the proliferation of AutoML. Whether it is possible to determine the effects of AutoML on AI products is a question I hope to take up in future research.

However, to conduct that research I will need to overcome another limitation of this study – my lack of technical knowledge about AI and machine learning. This dissertation

no doubt makes more than one misstep in theoretical or technical explication. While I made all possible attempts at accuracy, the nascent and continually evolving nature of the field makes it difficult. Going forward, I intend to pursue a fuller technical understanding of machine learning.

This dissertation was not only about characterizing work in the AI Industry. There are also limitations pertaining to the theoretical argument advanced on the basis of that analysis. The primary limitation here concerns generalizability. While AI work is certainly what *post-operaismo* describes as immaterial labour, it is not the only type of work that falls under this banner. It is not immediately certain that the critique advanced here applies to other types of immaterial labour such as, for instance, independent YouTube content production or other activities which do not occur in traditional workplaces. Other labour processes would have to be analyzed and attempts at automating them would have to be assessed. Nonetheless, given the evident importance of AI development in contemporary capitalism, the problems it poses for *post-operaismo*, which enjoys widespread popularity, are, I suggest, significant.

Finally, this dissertation might be accused of an overdriving pessimism or at least of lacking in constructive content by which critical research and activism concerning AI might proceed. While the argument advanced here is certainly a critique of the resolute optimism of *post-operaismo*, it does not intend to argue for a determinism in which the destiny of humanity is subjugation or annihilation at the hands of apotheotic cybernetic capital. I do agree with venture capitalist Kai-Fu Lee (2018), former executive at both Google and Microsoft, when he asserts that “if left unchecked, AI will throw gasoline on the socioeconomic fires” (161). The reader might have noticed that so far, this dissertation has not discussed at all how AI might be, as Lee puts it, “checked”. That is because somehow “checking” AI means very little if it does not include consideration of how AI is fundamentally tied up with the recursive technological process that is capital. In other words, “checking” AI means no less than “checking” capital’s inherent drive towards an increasing organic composition. While this is a task beyond the scope of this dissertation, now that the connections between AI and capital have been elucidated, I

will, in closing, briefly consider how critical thought and practice concerning the AI Industry might move forward.

### 7.3 Seize the Means of Cognition?

High-skill AI workers like data scientists and machine learning engineers already occupy positions of considerable power within cybernetic capitalism. If they follow up on the early stirrings of labour activism and organizing in tech work which were discussed in Chapter 4, it is impossible to predict what might happen. Indeed, it might be argued that those agitations are signs of the emerging autonomy which I have argued against in this dissertation. That is possible, though at this point all victories of labour against AI-producing capital have been decidedly local, temporary or marginal. While, for instance, after employee protests, Google dropped drone vision work on the Pentagon's Project Maven, it still insists that it will work with the military "in many other areas" (Statt and Vincent 2018). And while gender pay gaps may soon have to be openly disclosed by companies, over the AI Industry as a whole, the product of which is automation technology, "hangs the shadow that resistance may spur further automation" (Dyer-Witford, Kjösen and Steinhoff 2019, 108).

However, as I have argued briefly in Chapter 4, and elsewhere, AI is being positioned by capital as a new layer of infrastructure (Steinhoff 2019) or "means of cognition" (Dyer-Witford, Kjösen and Steinhoff 2019, 52). If AI does become a ubiquitous infrastructural technology, then the positions of high-skill AI workers within capital could become pivotal. Workers who service and control infrastructure can wield considerable social power by obstructing key moments of the valorization process. For instance, more than one analysis has pointed to the example of dock and port workers and their occupation of a "strategic position at the choke points of commodity distribution" within logistical capitalism to demonstrate how small sections of infrastructural workers can possess leverage against capital through their "ability to shut down large parts of the economy" (Bernes 2013; see also Frase 2015). If AI takes on the infrastructural ubiquity expected by its capitalist boosters, then AI workers will quite literally hold the controls to capital, at least until efforts to fully automate or "democratize" the complete process of AI production succeed.



Another possibility is suggested by what Gray and Suri (2019) call the paradox of automation's last mile or "the ever-moving frontier between what machines can and can't solve" (206). As we have seen, ghost workers are employed to complete key tasks in AI production that machines, for the moment, cannot. As data scientists figure out how to automate their tasks and crucial ghost workers become redundant, they are shuttled on to the next currently un-automatable task. AutoML techniques are attempting to automate ghost work entirely. However, during that window before they are replaced, ghost workers are critical to the functioning of AI systems and are therefore essential to the valorization processes of the capitals that employ them. Although their potential leverage is mitigated by their precarious employment situation and easy replaceability, as long as AI relies on ghost work, there will exist a weak point in cybernetic capital at the bottom, as well as the top, of the AI Industry labour hierarchy.

One might argue that expecting only the workers in the AI Industry to do something about AI is misguided. Perhaps an actual democratization of AI might occur. In other words, could AI be socially repurposed towards ends other than those of capital? This is the question posed, originally concerning capital's logistics networks, in the "reconfiguration debate" initiated by Toscano (2011) and Bernes (2013). I have previously argued that, currently, an anti-capitalist reconfiguration of AI seems implausible due to capital's extant control over the massive quantity of hardware, energy, data and skill required to produce and maintain such systems and the radical left's paucity of such capacities (Steinhoff 2017, 3-4). More recently, my co-authors and myself have pointed out that "the real subsumption of labour by capital means that capital develops and adopts technologies that fit its systemic requirements of valorization; this imperative can be baked into the very design of technology" (Dyer-Witthford, Kjösen and Steinhoff 2019, 149). While AI might be a necessity for a planned communist economy, as Cockshott (2017) argues, its utility to capital, which already controls it, is superlative. Control of AI by labour, were it possible, is unlikely to be wrested without a fight. Assessing the possibilities of a seizure of AI in the way that the Bolsheviks seized key infrastructure like railroads after the revolution in 1917, is, however, beyond the scope of this dissertation.

In closing, I want to suggest the possibility that AI might not necessarily have to be collectively seized to provide openings for critical action. It is possible that the technology of AI itself may afford openings. Zittrain (2019) suggest that using machine learning to solve problems incurs “intellectual debt” because “most machine-learning models cannot offer reasons for their ongoing judgments [and] there is no way to tell when they’ve misfired if one doesn’t already have an independent judgment about the answers they provide”. When machine learning applications are trained on data generated by other applications and AI systems are autonomously interacting with other AI systems, this debt grows. Zittrain (2019) argues that the widespread use of machine learning to solve problems and make decisions could result in a creation of a “world of knowledge without understanding ... a world without discernible cause and effect” in which humans rely on intelligent machines to interpret reality and guide them through an incomprehensible world.

However, it is not just the end user who experiences intellectual debt in this scenario. Zittrain’s understated, but most compelling, point is that the companies who produce AI will bear the brunt of intellectual debt. That is, their business processes will increasingly be based upon occult machine logics. The question is whether such an intellectual debt of capital could be exploited by anti-capitalist action. As Zittrain (2019) notes, machine learning misfires can “be triggered intentionally by someone who knows just what kind of data to feed into that system”. Research to counteract this looming problem is already underway. DARPA and others are invested in “explainable AI” research which seeks to mitigate or eliminate intellectual debt by building AI which can recount the processes which led to its output (Gunning 2017; Holzinger 2018). In DARPA’s case, this is presumably so that future autonomous weapons can explain why precisely they killed the targets that they did.

However, the AI world is not universally agreed on the desirability or feasibility of explainable AI. Deep learning luminary Hinton recently said in a *Wired* interview that he thinks it would be a “complete disaster” if regulators could “insist that you can explain how your AI system works” because even humans “can’t explain how they work, for most of the things they do” (quoted in Simonite 2018). A group of researchers in the field

assert that real solutions to the “explainable AI problem will be only made possible by truly interdisciplinary research, bridging data science and AI with human sciences, including philosophy and cognitive psychology” (Guidotti, Monreale, Pedreschi 2019). While the future of explainable AI is uncertain, the tendency towards automation without codification indicated by AutoML suggests a future rife with intellectual debt.

If capital’s increasing organic composition is achieved through automation without codification then its functioning will become increasingly mysterious even to itself. Companies will not know how or why they do what they do – only whether their capital is being valorized or not. The global system of capital will be a slew of such capitals interacting, all potentially with no knowledge of their own or anyone else’s methods. Land (2012) is correct when he describes capital as “not an essence but a tendency, the formula of which is decoding, or market-driven immanentization, progressively subordinating social reproduction to techno-commercial replication” (339-340). What he does not account for is that as capital migrates away from its historical human substratum it opens up new technological points of weakness. This suggests a way forward for critical thought and practice which cannot be explored here, but might be founded on a reconsideration of the young Negri’s (1979) advocacy for the working class sabotage of capitalist machines. Land (2012) suggests that “[o]nly proto-capitalism has ever been critiqued” (340). It is also true that only proto-capitalism has ever been attacked.

## References

All URLs, unless otherwise noted, were last visited on August 23, 2019.

- ABC 7 News. 2018. "Coolest employee perks at Silicon Valley tech companies." *ABC7 News*, July 24, 2018. <https://abc7news.com/technology/coolest-employee-perks-at-silicon-valley-tech-companies/3816443/>
- Acemoglu, Daron and Pascual Restrepo. 2018. "Artificial Intelligence, Automation and Work." National Bureau of Economic Research Working Paper #24196.
- Adam, Alison. 2006. *Artificial Knowing: Gender and the thinking machine*. Routledge.
- Adelson-Velskii, G.M., et al. 1970. "Programming a computer how to play chess." *Russian Mathematical Survey* 25 no. 221.
- Aghabozorghi, Saeed and Polong Lin. 2016. "Data Scientist vs Data Engineer, What's the difference?" *Cognitiveclass.ai*. June 6, 2016. <https://cognitiveclass.ai/blog/data-scientist-vs-data-engineer/>
- Agile Alliance. n.d. "Agile 101." Accessed August 23, 2019. <https://www.agilealliance.org/agile101/>
- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2016. "The simple economics of machine intelligence." *Harvard Business Review* 17.
- Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press: Boston.
- AI4ALL. n.d. "Our Story". <http://ai-4-all.org/about/our-story/>
- AI Index. 2017. "The AI Index 2017 Annual Report." <https://www.aiindex.org/2017-report.pdf>
- AI Index. 2018. "The AI Index 2018 Annual Report." <http://cdn.aiindex.org/2018/AI%20Index%202018%20Annual%20Report.pdf>
- AI Now. n.d. "Home Page." <https://ainowinstitute.org/>
- Ainspain, Nathan David. 1999. "The Geek shall inherit or Leave the Money and Run? Role identities and Turnover Decisions among Software Programmers and the High-Technology employees". PhD Diss, Cornell University.
- Allen, Kate. 2015. "How a Toronto professor's research revolutionized artificial intelligence." *The Star*. April 17, 2015. <https://www.thestar.com/news/world/2015/04/17/how-a-toronto-professors-research-revolutionized-artificial-intelligence.html>

- Allen, Gregory C. 2019. "Understanding China's AI Strategy." *Center for a New American Security*. February 6. <https://www.cnas.org/publications/reports/understanding-chinas-ai-strategy>
- Alpaydin, Ethem. 2014. *Introduction to Machine Learning*. MIT Press.
- Alpaydin, Ethem. 2016. *Machine Learning: the New AI*. MIT Press.
- Alquati, Romano. 2013 [1961]. "Organic Composition of Capital and Labour-Power at Olivetti." *Viewpoint Magazine* 27. September 2013. Accessed 25 April 2018. <https://www.viewpointmag.com/2013/09/27/organic-composition-of-capital-and-labor-power-at-olivetti-1961/>
- Amadeo, Ron. 2018. "Google's iron grip on Android: Controlling open source by any means necessary." *Ars Technica*. July 21, 2018. <https://arstechnica.com/gadgets/2018/07/googles-iron-grip-on-android-controlling-open-source-by-any-means-necessary/>
- Amin, Ash. 1994. "Post-Fordism: models, fantasies and phantoms of transition." In *Post-Fordism: a reader*, edited by Ash Amin. Wiley-Blackwell.
- Amodei, Dario and Danny Hernandez. 2018. "AI and Compute." *OpenAI Blog*. May 16, 2018. <https://openai.com/blog/ai-and-compute/>
- Anderson, Jess. 2018. "Data engineers vs data scientists." *O'Reilly*. April 11, 2018. <https://www.oreilly.com/ideas/data-engineers-vs-data-scientists>
- Anderson, Michael, and Susan Leigh Anderson, eds. 2011. *Machine Ethics*. Cambridge University Press
- Andrews, Chris K., Craig D. Lair and Bart Landry. 2005. "The labor process in software startups: production on a virtual assembly line?" In *Management, Labour Process and Software Development: Reality Bites*, edited by Rowena Barret. Routledge.
- Angeli, C. 2010. "Diagnostic expert systems: From expert's knowledge to real-time systems." *Advanced knowledge based systems: Model, applications & research* 1: 50-73.
- Angwin, Julia, Jeff Larson, Surya Mattu and Lauren Kirchner. 2016 "Machine Bias." *ProPublica*. May 23, 2016. Accessed November 7, 2018. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Arkin, Ronald. 2009. *Governing lethal behavior in autonomous robots*. CRC Press.
- Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. 2016. "The risk of automation for jobs in OECD countries: A comparative analysis." OECD Social, Employment, and Migration Working Papers 189.
- Artificial Intelligence Committee. 2017. "AI in the UK: Ready, willing, and able? Report of Session 2017-19." *House of Lords*. April 16, 2017. <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/10002.htm>

- Athanasiou, Tom. 1985. "Artificial Intelligence: cleverly disguised politics." In *Compulsive Technology: Computers as Culture*, edited by Tony Solomonides and Les Levidow, 13-35. London: Free Association Books.
- Babbage, Charles. 1832. *On The Economy of Machinery and Manufactures*. <https://www.fulltextarchive.com/pdfs/On-the-Economy-of-Machinery-and-Manufactures.pdf>
- Bardini, Thierry. 2000. *Bootstrapping: Douglas Engelbart, coevolution, and the origins of personal computing*. Stanford University Press.
- Barrat, James. 2013. *Our final invention: Artificial intelligence and the end of the human era*. Macmillan.
- Barret, Brian. 2017. "What Amazon Echo and Google Home do with your voice data." *Wired*. November 24, 2017. <https://www.wired.com/story/amazon-echo-and-google-home-voice-data-delete/>
- Barrett, Rowena. 2001. "Labouring under an illusion? The labour process of software development in the Australian information industry." *New Technology, Work and Employment* 16 (1): 18-34.
- Barrett, Rowena. 2004. "Working at Webboyz: An analysis of control over the software development labour process." *Sociology* 38 (4): 777-794.
- Barrett, Rowena. 2005. "Managing the software development labour process: direct control, time and technical autonomy." In *Management, Labour Process and Software Development: Reality Bites*, edited by Rowena Barret. Routledge.
- Bastani, Aaron. 2019. *Fully Automated Luxury Communism*. Verso Books.
- Baum, Seth. 2017. "A Survey of Artificial General Intelligence Projects for Ethics, Risk, and Policy." Global Catastrophic Risk Institute Working Paper 17-1. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3070741](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3070741)
- Beck, Kent, Mike Beedle, Arie Van Bennekum, Alistair Cockburn, Ward Cunningham, Martin Fowler, James Grenning, Jim Highsmith, Andrew Hunt, Ron Jeffries, Jon Kern, Brian Marick, Robert C. Martin, Steve Mellor, Ken Schwaber, Jeff Sutherland and Dave Thomas. 2001. "Manifesto for agile software development." <http://agilemanifesto.org/>
- Beer, Randall D. 2014. "Dynamical Systems and Embedded Cognition." In *The Cambridge Handbook of Artificial Intelligence*, edited by Keith Frankish and William M. Ramsey, 128-48. Cambridge: Cambridge University Press.
- Bell, Daniel. 1973. *The Coming of Post-industrial Society. A Venture in Social Forecasting*. New York: Basic Books.
- Benaich, Nathan. 2016. "State of the AI Market in 2015: A focus on investments." Playfair Capital. [https://docs.google.com/presentation/d/1R59\\_8WH-TN\\_rQOTSRnQdgDusiJM5COpKJ7mls-2FQvo/edit#!](https://docs.google.com/presentation/d/1R59_8WH-TN_rQOTSRnQdgDusiJM5COpKJ7mls-2FQvo/edit#!)
- Benner, Katie. 2017a. "Women in Tech Speak Frankly on Culture of Harassment." *The New York Times*. June 30, 2017.

<https://www.nytimes.com/2017/06/30/technology/women-entrepreneurs-speak-out-sexual-harassment.html?hpw=&rref=technology&action=click&pgtype=Homepage&module=inline&region=bottom-well&WT.nav=bottom-well>

- Benner, Katie. 2017b. "A Backlash Builds Against Sexual Harassment in Silicon Valley." *The New York Times*. July 3, 2017.  
<https://www.nytimes.com/2017/07/03/technology/silicon-valley-sexual-harassment.html?action=click&module=RelatedCoverage&pgtype=Article&region=Footer>
- Berardi, Franco. 2009. *The soul at work: From alienation to autonomy*. Semiotext (e).
- Berardi, Franco. 2011. *After the Future*. AK Press.
- Berg, Janine, Marianne Furrer, Ellie Harmon, Uma Rani and M Six Silberman. 2018. "Digital Labour Platforms and the Future of Work: Towards Decent Work in the Online World". *International Labour Organisation*.  
[www.ilo.org/wcmsp5/groups/public/---dgreports/---dcomm/---publ/documents/publication/wcms\\_645337.pdf](http://www.ilo.org/wcmsp5/groups/public/---dgreports/---dcomm/---publ/documents/publication/wcms_645337.pdf)
- Bergen, Mark and Kurt Wagner. 2015. "Welcome to the AI Conspiracy: The 'Canadian Mafia' behind tech's latest craze." *Vox*. July 15, 2015.  
<https://www.vox.com/2015/7/15/11614684/ai-conspiracy-the-scientists-behind-deep-learning>
- Bergvall-Kåreborn, Birgitta, and Debra Howcroft. 2013. "The future's bright, the future's mobile": a study of Apple and Google mobile application developers." *Work, employment and society* 27 (6): 964-981.
- Berman, Bruce. 1992. "Artificial Intelligence and the Ideology of Capitalist Reconstruction." *AI & Society* 6, 103-114.
- Bernes, Jasper. 2013. "Logistics, Counterlogistics and the Communist Prospect." *Endnotes* 3. <https://endnotes.org.uk/issues/3/en/jasper-bernes-logisticscounterlogistics-and-the-communist-prospect>
- Bhasin, Sanjay, and Peter Burcher. 2006. "Lean viewed as a philosophy." *Journal of Manufacturing Technology Management* 17 (1): 56-72.
- Bloomberg News. 2017. "China's Got A Huge Artificial Intelligence Plan." Bloomberg. July 20, 2017. <https://www.bloomberg.com/news/articles/2017-07-21/china-artificial-intelligence-plan-seeks-59-billion-industry>
- Boden, Margaret A., ed. 1990. *The Philosophy of Artificial Intelligence*. New York: Oxford University Press.
- Boden, Margaret A. 2014. "GOF AI." In *The Cambridge Handbook of Artificial Intelligence*, edited by Keith Frankish and William M. Ramsey, 89-107. Cambridge: Cambridge University Press.
- Boden, Margaret A. 2016. *AI: Its Nature and Future*. Oxford, UK: Oxford University Press.

- Bonefeld, Werner. 2014. *Critical theory and the critique of political economy: On subversion and negative reason*. Bloomsbury Publishing USA.
- Bonnefon, Jean-François, Azim Shariff, and Iyad Rahwan. 2016. "The social dilemma of autonomous vehicles." *Science* 352 (6293): 1573-1576.
- Bostrom, Nick. 2001. "What is Transhumanism?"  
<https://nickbostrom.com/old/transhumanism.html>
- Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.
- Bowne-Anderson, Hüge. 2018. "What Data Scientists Really Do, According to 35 Data Scientists." *Harvard Business Review*. August 15, 2018. [https://hbr.org/2018/08/what-data-scientists-really-do-according-to-35-data-scientists?referral=03758&cm\\_vc=rr\\_item\\_page.top\\_right](https://hbr.org/2018/08/what-data-scientists-really-do-according-to-35-data-scientists?referral=03758&cm_vc=rr_item_page.top_right)
- Boyle, Alan. 2017. "IBM makes a 10-year, \$240M investment in artificial intelligence research at MIT". *GeekWire*. September 6, 2017. <https://www.geekwire.com/2017/ibm-makes-240m-investment-artificial-intelligence-research-mit/>
- Bratton, Benjamin H. 2016. *The Stack: On Software and Sovereignty*. MIT Press.
- Braidotti, Rosi. 2013. *The Posthuman*. Cambridge: Polity.
- Brassier, Ray. 2014. "Prometheanism and its Critics." In *#Accelerate: The Accelerationist Reader*, edited by Robin Mackay and Arven Avanesian, 486-87. Urbanomic.
- Braverman, Harry. 1998. *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. NYU Press.
- Briken, Kendra, Shiona Chillas, and Martin Krzywdzinski. 2017. *The new digital workplace: How new technologies revolutionise work*. Macmillan International Higher Education.
- Brinkman, Svend. 2013. *Qualitative Interviewing*. Oxford University Press.
- Brockman, Greg. 2019. "Microsoft Invest In and Partners With OpenAI To Support Us Building Beneficial AI." *OpenAI Blog*. July 22, 2019. <https://openai.com/blog/microsoft/>
- Brooks, Rodney A. 1991. "Intelligence without representation." *Artificial intelligence* 47 (1-3): 139-159.
- Brown, Brian A. 2012. "Will work for free: Examining the biopolitics of unwaged immaterial labour." PhD diss., University of Western Ontario.
- Brownlee, Jason. 2013. "How to Prepare Data For Machine Learning." *Machine Learning Mastery*. December 25, 2013. <https://machinelearningmastery.com/how-to-prepare-data-for-machine-learning/>
- Brownlee, Jason. 2014. "Discover Feature Engineering, How to Engineer Features and How to Get Good at It." *Machine Learning Mastery*. September 26,



2014. <https://machinelearningmastery.com/discover-feature-engineering-how-to-engineer-features-and-how-to-get-good-at-it/>
- Brownlee, Jason. 2015. "Gentle Introduction to Predictive Modeling." *Machine Learning Mastery*. September 20, 2015. <https://machinelearningmastery.com/gentle-introduction-to-predictive-modeling/>
- Brownlee, Jason. 2016. "Supervised and Unsupervised Machine Learning Algorithms." *Machine Learning Mastery*. March 16, 2016. <https://machinelearningmastery.com/supervised-and-unsupervised-machine-learning-algorithms/>
- Brundage, Miles, Shahar Avin, Jack Clark, Helen Toner, Peter Eckersley, Ben Garfinkel, Allan Dafoe, Paul Scharre, Thomas Zeitzoff, Bobby Filar, Hyrum Anderson, Heather Roff, Gregory C. Allen, Jacob Steinhardt, Carrick Flynn, Seán Ó hÉigeartaigh, Simon Beard, Haydn Belfield, Sebastian Farquhar, Clare Lyle, Rebecca Crootof, Owain Evans, Michael Page, Joanna Bryson, Roman Yampolskiy and Dario Amodei. 2018. *The Malicious Use of Artificial Intelligence: Forecasting, Prevention and Mitigation*. [https://img1.wsimg.com/blobby/go/3d82daa4-97fe-4096-9c6b-376b92c619de/downloads/1c6q2kc4v\\_50335.pdf](https://img1.wsimg.com/blobby/go/3d82daa4-97fe-4096-9c6b-376b92c619de/downloads/1c6q2kc4v_50335.pdf)
- Bruno, Alessandro. 2018. "Weak Post-2008 Crisis Economic Recovery Puts Capitalism at Risk." *Lombardi Letter*. April 5, 2018. <https://www.lombardiletter.com/weak-post-2008-crisis-economic-recovery-puts-capitalism-at-risk/27286/>
- Brynjolfsson, Erik, and Andrew McAfee. 2014. *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.
- Brynjolfsson, Erik, and Andrew McAfee. 2017. "The business of Artificial Intelligence: What it can--and cannot--do for your organization." *Harvard Business Review Digital Articles* 7: 3-11.
- Buhr, Sarah. 2017. "Tech employees protest in front of Palantir HQ over fears it will build Trump's Muslim registry." *TechCrunch*. January 18, 2017. <https://techcrunch.com/2017/01/18/tech-employees-protest-in-front-of-palantir-hq-over-fears-it-will-build-trumps-muslim-registry/>
- Burston, Jonathan, Nick Dyer-Witford, and Alison Hearn. 2010. "Digital labour: Workers, authors, citizens." *Special issue, ephemera: theory and politics in organization* 10 (3/4): 214-221.
- Caffentzis, George. 2013. *In letters of blood and fire: Work, machines, and the crisis of capitalism*. pm Press.
- Cameron, Rondo. 1982. "The industrial revolution: A misnomer." *The History Teacher* 15 (3): 377-384.
- Camfield, David. 2007. "The multitude and the kangaroo: A critique of Hardt and Negri's theory of immaterial labour." *Historical Materialism* 15 (2): 21-52.

- Campbell, Virginia. 2004. "How RAND invented the Postwar World." <http://www.rand.org/content/dam/rand/www/external/about/history/Rand.IT.Summer04.pdf>
- Campbell-Kelly, Martin. 2004. *From airline reservations to Sonic the Hedgehog: a history of the software industry*. MIT Press.
- Canadian Federal Government. 2017. "Building a Strong Middle Class (Budget 2017)." <https://www.budget.gc.ca/2017/docs/plan/budget-2017-en.pdf>
- Canon, Gabrielle. 2018. "Google gave top executive \$90m payoff but kept sexual misconduct claim quiet – report." *The Guardian*. October 25, 2018. <https://www.theguardian.com/technology/2018/oct/25/google-andy-rubin-android-creator-payoff-sexual-misconduct-report>
- Carr, Nicholas. 2014. *The Glass Cage: Automation and Us*. WW Norton & Company.
- Case, Nicky. 2018. "How To Become A Centaur." *Journal of Design and Science*. February 6, 2018. <https://jods.mitpress.mit.edu/pub/issue3-case>
- Castells, Manuel, and Yuko Aoyama. 1994. "Paths towards the informational society: Employment structure in G-7 countries, 1920-90." *International Labour Review* 133.
- Cath, Corinne, Sandra Wachter, Brent Mittelstadt, Mariarosaria Taddeo, and Luciano Floridi. 2018. "Artificial intelligence and the 'good society': the US, EU, and UK approach." *Science and Engineering Ethics* 24 (2): 505-528.
- Chang, Emily. 2018. "'Oh My God, This Is So F---ed UP': Inside Silicon Valley's Secretive, Orgiastic Dark Side." *Vanity Fair*. January 2, 2018. <https://www.vanityfair.com/news/2018/01/brotopia-silicon-valley-secretive-orgiastic-inner-sanctum>
- Chang, Emily. 2019. *Brotopia: Breaking Up the Boys' Club of Silicon Valley*. Portfolio.
- Chen, Angela. 2019. "How to regulate Big Tech without breaking it up". *MIT Technology Review*. June 7, 2019. <https://www.technologyreview.com/s/613640/big-tech-monopoly-breakup-amazon-apple-facebook-google-regulation-policy/>
- Chen, Frank. 2016. "AI, Deep Learning, and Machine Learning: A Primer," *Andreessen Horowitz*. June 10, 2016. <http://a16z.com/2016/06/10/ai-deep-learning-machines>
- Chen, Wenjie, Mico Mrkaic and Malhar Nabar. 2018. "Lasting Effects: The Global Economic Recovery 10 Years After The Crisis". *IMF Blog*. October 3, 2018. <https://blogs.imf.org/2018/10/03/lasting-effects-the-global-economic-recovery-10-years-after-the-crisis/>
- Chun, Wendy Hui Kyong. 2005. "On software, or the persistence of visual knowledge." *grey room* 18: 26-51.
- Churchland, Patricia Smith. 1989. *Neurophilosophy: Toward a unified science of the mind-brain*. MIT Press.

- Churchland, Patricia Smith, and Terrence J. Sejnowski. 1994. *The Computational Brain*. MIT Press.
- Churchland, Paul M. 2013. *Matter and Consciousness*. MIT Press.
- CIO DoD. 2018. "Joint Artificial Intelligence Center." Chief Information Officer U.S. Department of Defense. <https://dodcio.defense.gov/About-DoD-CIO/Organization/JAIC/>
- Cleaver, Harry. 1992. "The Inversion of Class Perspective in Marxian Theory: From Valorisation to Self-Valorisation." In *Open Marxism Volume 2: Theory and Practice*, edited by Werner Bonefeld, Richard Gunn and Kosmas Psychopedis. Pluto Press.
- Cleaver, Harry. 2000 [1979]. *Reading Capital Politically*. AntiTheses and AK Press.
- Cockshott, Paul. 2017. "Big Data and Supercomputers: Foundations of Cyber Communism." <http://paulcockshott.co.uk/publication-archive/Talks/politiceconomy/Big%20Data%20and%20Super-Computers%20foundations%20of%20Cyber%20Communism.pdf>
- Cohn, Mike. 2010. *Succeeding with Agile: software development using Scrum*. Pearson Education.
- Colson, Eric. 2019. "Why Data Science Teams Need Generalists, Not Specialists." *Harvard Business Review*. March 8, 2019. <https://hbr.org/2019/03/why-data-science-teams-need-generalists-not-specialists>
- Conger, Kate and Noam Scheiber. 2019. "Employee activism is alive in tech. It stops short of organising unions." *The New York Times*. July 8, 2019. <https://www.nytimes.com/2019/07/08/technology/tech-companies-union-organizing.html>
- Cook, James. 2018. "Amazon scraps 'sexist AI' recruiting tool that showed bias against women." *The Telegraph*. October 10, 2018. <https://www.telegraph.co.uk/technology/2018/10/10/amazon-scraps-sexist-ai-recruiting-tool-showed-bias-against/>
- Cook, Tim. 2019. "You Deserve Privacy Online. Here's How You Could Actually Get It." *Time*. January 16, 2019. <https://time.com/collection-post/5502591/tim-cook-data-privacy/>
- Cooley, Mike. 1980. "Computerization: Taylor's Latest Disguise." *Economic and Industrial Democracy* 1 (4): 523-539.
- Cooley, Mike. 1981. "On The Taylorisation of Intellectual Work." In *Science, Technology and the Labour Process Volume 2*, edited by Les Levidow and Robert Young. CSE Books.
- Cooley, Mike. 1987. *Architect or Bee? The Human Price of Technology*. Hogarth Press.
- Copeland, Jack. 2000. "What is Artificial Intelligence?" *AlanTuring.net*. Accessed November 7,

2018. [http://www.alanturing.net/turing\\_archive/pages/reference%20articles/what\\_is\\_AI/What%20is%20AI11.html](http://www.alanturing.net/turing_archive/pages/reference%20articles/what_is_AI/What%20is%20AI11.html)
- Copeland, Jack. 1993. *Artificial Intelligence: A Philosophical Introduction*. Blackwell Publishers.
- Corbet, Jonathan and Greg Kroah-Hartman. 2018. "2017 Linux Kernel Development Report." *Linux Foundation*. <https://www.linuxfoundation.org/2017-linux-kernel-report-landing-page/>
- Coren, Michael J. 2017. "Silicon valley workers are talking about starting their first union in 2017 to resist Trump." *Quartz*. March 24, 2017. Accessed 28 June 2018. <https://qz.com/916534/silicon-valley-tech-workers-are-talking-about-starting-their-first-union-in-2017-to-resist-trump/>
- Costine, Josh. 2019. "Facebook announces Libra cryptocurrency: All you need to know." *TechCrunch*. June 18, 2019. <https://techcrunch.com/2019/06/18/facebook-libra/>
- Crawford, Kate. 2016. "Artificial Intelligence's White Guy Problem." *The New York Times*. June 26, 2016. <https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html>
- Crevier, Daniel. 1993. *AI: the tumultuous history of the search for artificial intelligence*. Basic Books.
- Crowdfunder. 2016. "2016 Data Science Report." [https://visit.figure-eight.com/rs/416-ZBE-142/images/CrowdFlower\\_DataScienceReport\\_2016.pdf](https://visit.figure-eight.com/rs/416-ZBE-142/images/CrowdFlower_DataScienceReport_2016.pdf)
- Crowe, Steve. 2018. "Rethink Robotics IP acquired by HAHN Group." *The Robot Report*. October 25, 2018. <https://www.therobotreport.com/hahn-group-acquires-rethink-robotics-ip/>
- Culurciello, Eugenio. 2017. "Navigating the Unsupervised Learning Landscape." *Medium*. May 4, 2017. <https://medium.com/intuitionmachine/navigating-the-unsupervised-learning-landscape-951bd5842df9>
- Cummings, M.L. 2018. "Artificial Intelligence and the Future of Warfare." In *Artificial Intelligence and International Affairs: Disruption Anticipated*, 7-18. Chatham House.
- Cummings, M.L., Heather M. Roff, Kenneth Cukier, Jacob Parakilas and Hannah Bryce. 2018. *Artificial Intelligence and International Affairs: Disruption Anticipated*. Chatham House. <https://www.chathamhouse.org/sites/default/files/publications/research/2018-06-14-artificial-intelligence-international-affairs-cummings-roff-cukier-parakilas-bryce.pdf>
- Danaher, John. 2019. "The rise of the robots and the crisis of moral patiency." *AI & SOCIETY* 34 (1): 129-136.
- Darrow, Barb. 2017. "Microsoft Bids Goodbye to 'Mobile First' Mantra in Favor of AI." *Fortune*. August 3, 2017. <https://fortune.com/2017/08/03/microsoft-cloud-ai-mobile/>

- Dash, Anil. 2012. "The Blue Collar Coder." *Anil Dash's Blog*. October 5. [https://anildash.com/2012/10/05/the\\_blue\\_collar\\_coder/](https://anildash.com/2012/10/05/the_blue_collar_coder/)
- DataLoop. n.d. "Machine Learning Ready-to-Use Datasets." <https://www.dataloop.ai/create-your-datasets#!>
- DataRobot. n.d. "Machine Learning Model Deployment." <https://www.datarobot.com/wiki/machine-learning-model-deployment/>
- Daubs, Michael S., and Vincent R. Manzerolle, eds. 2017. *Mobile and Ubiquitous Media: Critical and International Perspectives*. Peter Lang International Academic Publishers.
- Daugherty, Paul R., and H. James Wilson. 2018. *Human+ Machine: Reimagining Work in the Age of AI*. Harvard Business Review Press.
- Davenport, Thomas H. 2005. *Thinking for a Living: how to get better performances and results from knowledge workers*. Harvard Business Press.
- Davenport, Thomas H., and Julia Kirby. 2016. *Only humans need apply: winners and losers in the age of smart machines*. Harper Business.
- Davenport, Thomas H. and D.J. Patil. 2012. "Data Scientist: The Sexiest Job of the 21<sup>st</sup> Century." *Harvard Business Review*. October 2012. Accessed 13 September 2018. <https://hbr.org/2012/10/data-scientist-the-sexiest-job-of-the-21st-century>
- Davis, Jessica. 2016. "7 Data and Analytics Jobs to Know." *Information Week*. October 25, 2016. Accessed 13 September 2018. [https://www.informationweek.com/big-data/7-data-and-analytics-job-titles-to-know--/d/d-id/1327272?image\\_number=8](https://www.informationweek.com/big-data/7-data-and-analytics-job-titles-to-know--/d/d-id/1327272?image_number=8)
- Davis, Randall, Howard Austin, Ingrid Carlbom, Bud Frawley, Paul Pruchnik, Rich Sneiderman, and J. A. Gilreath. 1981. "The DIPMETER ADVISOR: Interpretation of Geologic Signals." In *Proceedings of the 7<sup>th</sup> International Joint Conference on Artificial Intelligence Volume 2*: 846-849. Morgan Kaufmann Publishers Inc.
- Day, Matt. 2019. "Google, Amazon Expanding Smart Home Device Data Gathering." *Insurance Journal*. February 12, 2019. <https://www.insurancejournal.com/news/national/2019/02/12/517456.htm>
- DeepMind. n.d.a. "Home Page." <https://deepmind.com/>
- DeepMind. n.d.b. "About." <https://deepmind.com/about/>
- DeepMind. n.d.c. "DeepMind Health." <https://deepmind.com/applied/deepmind-health/>
- Deleuze, Gilles and Félix Guattari. 1983. *Anti-Oedipus: Capitalism and Schizophrenia*, translated by Robert Hurley, Mark Seem, and Helen R. Lane. Minneapolis: University of Minnesota Press.
- Deleuze, Gilles, and Félix Guattari. 1987. *A Thousand Plateaus: Capitalism and Schizophrenia*, translated by Brian Massumi. Minneapolis, University of Minnesota Press.

- Dennett, D. C. 2006. "Cognitive Wheels: The frame problem of AI." In *Philosophy of Psychology: Contemporary Readings*, edited by José Luis Bermúdez. New York: Routledge.
- Denning, Steve. 2015. "Agile: The World's Most Popular Innovation Engine." *Forbes*. July 23, 2015. <https://www.forbes.com/sites/stevedenning/2015/07/23/the-worlds-most-popular-innovation-engine/#77ccfe27c769>
- De Stefano, Valerio. 2015. "The rise of the just-in-time workforce: On-demand work, crowdwork, and labor protection in the gig-economy." *Comparative Labor Law and Policy Journal* 37(3): 471-504.
- Dettmers, Tim. 2015. "Deep Learning in a Nutshell: Core Concepts." *NVIDIA Developer Blog*. November 3, 2015. Accessed November 5, 2018. <https://devblogs.nvidia.com/deep-learning-nutshell-core-concepts/>
- De Vries, Jan. 1994. "The industrial revolution and the industrious revolution." *The Journal of Economic History* 54 (2): 249-270.
- Dickson, Ben. 2018. "Artificial Intelligence Has a Bias Problem, and It's Our Fault." *PC Magazine*. June 14, 2018. <https://www.pcmag.com/article/361661/artificial-intelligence-has-a-bias-problem-and-its-our-fau>
- Dickey, Megan Rose. 2014. "It's time to admit the amount of information Google gathers about us is terrifying." *Business Insider*. June 18, 2014. <https://www.businessinsider.com/the-information-google-is-gathering-about-us-is-terrifying-2014-6>
- Dinerstein, A. C., and Neary, M., eds. 2002. *The Labour Debate: An investigation into the theory and reality of capitalist work*. Aldershot: Ashgate.
- Ding, Jeffrey. 2018. "Deciphering China's AI Dream: The context, components, capabilities and consequences of China's strategy to lead the world in AI." *Future of Humanity Institute*. <https://www.fhi.ox.ac.uk/deciphering-chinas-ai-dream/>
- DiNucci, Darcy. 1999. "Fragmented future." *Print* 53 (4): 32-33.
- Doctorow, Cory. 2019. "Regulating Big Tech makes them stronger, so they need competition instead." *The Economist*. June 6, 2019. <https://www.economist.com/open-future/2019/06/06/regulating-big-tech-makes-them-stronger-so-they-need-competition-instead>
- Docherty, Bonnie. 2012. "Losing Humanity: The Case Against Killer Robots." *Human Rights Watch*. <https://www.hrw.org/report/2012/11/19/losing-humanity/case-against-killer-robots>
- Dolata, Ulrich. 2018. "Internet Companies: Market Concentration, Competition and Power." In *Collectivity and Power on the Internet: A Sociological Perspective*, edited by Ulrich Dolata and Jan-Felix Schrape. Springer.
- Domingos, Pedro. 2015. *The Master Algorithm: How the quest for the ultimate learning machine will remake our world*. Basic Books.

- D'Onfro, Jillian. 2016. "Google CEO: We're headed towards an 'AI first' world." *Business Insider India*. April 22, 2016. <https://www.businessinsider.in/Google-CEO-Were-headed-towards-an-AI-first-world/articleshow/51935048.cms>
- Dong, Catherine. 2017. "The evolution of machine learning." *TechCrunch*. August 18, 2017. Accessed 12 September 2018. <https://techcrunch.com/2017/08/08/the-evolution-of-machine-learning/>
- Dreyfus, Hubert L. 1979. *What computers can't do: The limits of artificial intelligence*. New York: Harper & Row.
- Dreyfus, Hubert L. 1992. *What computers still can't do: a critique of artificial reason*. MIT Press.
- Drucker, Peter. 1957. *Landmarks of Tomorrow*. New York: Harper and Brothers.
- Duke, Sue. 2018. "Will AI make the gender gap in the workplace harder to close?" *World Economic Forum*. December 21, 2018. <https://www.weforum.org/agenda/2018/12/artificial-intelligence-ai-gender-gap-workplace/>
- Dutton, Tim, Brent Barron and Gaga Boskovic. 2018. "Building an AI World: Report on National and Regional AI Strategies." *CIFAR*. [https://www.cifar.ca/docs/default-source/ai-society/buildinganaiworld\\_eng.pdf](https://www.cifar.ca/docs/default-source/ai-society/buildinganaiworld_eng.pdf)
- Dutton, Tim. 2018. "An Overview of National AI Strategies." *Medium*. June 28, 2018. <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd>
- Dyer-Witheford, Nick. 1999. *Cyber-Marx: Cycles and Circuits of Struggle in High-Technology Capitalism*. University of Illinois Press.
- Dyer-Witheford, Nick. 2001. "Empire, immaterial labor, the new combinations, and the global worker." *Rethinking Marxism* 13 (3-4): 70-80.
- Dyer-Witheford, Nick. 2005. "Cyber-Negri: General intellect and immaterial labor." In *The Philosophy of Antonio Negri Volume 1*, edited by Timothy Murphy and Abdul-Karim Mustapha. Pluto Press.
- Dyer-Witheford, Nick. 2013. "Red Plenty Platforms." *Culture Machine* 14.
- Dyer-Witheford, Nick. 2015. *Cyber-Proletariat: Global Labour in the Digital Vortex*. Pluto Press.
- Dyer-Witheford, Nick, and Svitlana Matviyenko. 2019. *Cyberwar and Revolution: Digital Subterfuge in Global Capitalism*. University of Minnesota Press.
- Dyer-Witheford, Nick, Atle Mikkola Kjøsén and James Steinhoff. 2019. *Inhuman Power: Artificial Intelligence and the Future of Capitalism*. Pluto Press.
- Eckert, Daniel. 2016. "AI is the new electricity." *Price Waterhouse Coopers Next in Tech Blog*. July 11, 2016. Accessed 31 August 2017. <http://usblogs.pwc.com/emerging-technology/ai-is-the-new-electricity/>

- Economist, The. 2016a. "The Return of the Machinery Question." June 23, 2016. <https://www.economist.com/special-report/2016/06/23/the-return-of-the-machinery-question>
- Economist, The. 2016b. "The rise of the superstars." September 15, 2016. <https://www.economist.com/special-report/2016/09/15/the-rise-of-the-superstars>
- Economist, The. 2017. "Google leads in the race to dominate artificial intelligence." December 7, 2017. <https://www.economist.com/business/2017/12/07/google-leads-in-the-race-to-dominate-artificial-intelligence>
- Economist, The. 2019. "Competition, not break-up, is the cure for tech giants' dominance." March 13, 2019. <https://www.economist.com/business/2019/03/13/competition-not-break-up-is-the-cure-for-tech-giants-dominance>
- Edwards, Paul N. 1996. *The Closed World: Computers and the Politics of Discourse in Cold War America*. Cambridge: MIT Press.
- Edwards, Jim. 2016. "One of the kings of the '90s dot-com bubble now faces 20 years in prison." *Business Insider*. December 6, 2016. <https://www.businessinsider.com/where-are-the-kings-of-the-1990s-dot-com-bubble-bust-2016-12>
- Ehrenkranz, Melanie. 2017. "Google Employees Organize Their Own Study of Gender Pay Gap." *Gizmodo*. August 9, 2017. <https://gizmodo.com/google-employees-organize-their-own-study-of-gender-pay-1802767010>
- Ekbia, Hamid Reza. 2008. *Artificial Dreams: The Quest for Non-Biological Intelligence*. Cambridge University Press.
- Elbe, Ingo. 2013. "Between Marx, Marxism, and Marxisms – Ways of Reading Marx's Theory." *Viewpoint Magazine*. October 21, 2013. <https://www.viewpointmag.com/2013/10/21/between-marx-marxism-and-marxisms-ways-of-reading-marxs-theory/>
- Elsken, Thomas, Jan Hendrik Metzen, and Frank Hutter. 2019. "Neural architecture search: A survey." *Journal of Machine Learning Research* 20 (55): 1-21.
- Emerj. n.d. "Companies." <https://emerj.com/companies/>
- Endnotes. 2010. "Misery and Debt: On the Logic and History of Surplus Populations and Surplus Capital." *Endnotes* 2. <https://endnotes.org.uk/issues/2/en/endnotes-misery-and-debt>
- Engels, Frederick. 1947. *Anti-Dühring*. Progress Publishers. <https://www.marxists.org/archive/marx/works/1877/anti-duhring/index.htm>
- Ensmenger, Nathan and William Aspray. 2002. "Software as labor process." In *History of Computing: Software Issues*, 139-165. Berlin: Springer.
- Ertel, Wolfgang. 2018. *Introduction to Artificial Intelligence*. Springer.



- Esteva, Andre, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. 2017. "Dermatologist-level classification of skin cancer with deep neural networks." *Nature* 542 (7639): 115.
- Etzioni, Oren. 2018. "AI advances require 99 percent manual labor." *MIT Technology Review*. March 26, 2018. <https://www.technologyreview.com/video/610647/ai-advances-require-99-percent-manual-labor/>
- Eubanks, Virginia. 2018. *Automating Inequality: How High-tech Tools Profile, Police and Punish the Poor*. St. Martin's Press.
- Evans, Richard and Jim Gao. 2016. "DeepMind AI Reduces Google Data Centre Cooling Bill by 40%." *DeepMind Blog*. July 20, 2016. <https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/>
- Executive Office of the President and National Science and Technology Council Committee on Technology. 2016. "Preparing for the Future of Artificial Intelligence." October. [https://obamawhitehouse.archives.gov/sites/default/files/whitehouse\\_files/microsites/ostp/NSTC/preparing\\_for\\_the\\_future\\_of\\_ai.pdf](https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf)
- Executive Office of the President. 2016. "Artificial Intelligence, Automation and the Economy." December. <https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.pdf>
- Fabian. 2018. "Global Artificial Intelligence Landscape." *Medium*. May 22, 2018. <https://medium.com/@bootstrappingme/global-artificial-intelligence-landscape-including-database-with-3-465-ai-companies-3bf01a175c5d>
- Faggella, Daniel. 2019. "Artificial Intelligence Industry – An Overview by Segment." *Emerj*. <https://emerj.com/ai-sector-overviews/artificial-intelligence-industry-an-overview-by-segment/>
- Fang, Lee. 2018. "Leaked emails show Google expected lucrative military drone AI work to grow exponentially." *The Intercept*. May 31, 2018. <https://theintercept.com/2018/05/31/google-leaked-emails-drone-ai-pentagon-lucrative/>
- Farivar, Cyrus. 2018. "In Iceland, Bitcoin mining will soon use more energy than its residents." *Ars Technica*. February 12, 2018. <https://arstechnica.com/tech-policy/2018/02/in-iceland-bitcoin-mining-will-soon-use-more-energy-than-its-residents/>
- Federal Reserve. 2014. "Changes in U.S. Family Finances from 2010 to 2013: Evidence from the Survey of Consumer Finances." *Federal Reserve Bulletin* 100 (4). September. <https://www.federalreserve.gov/pubs/bulletin/2014/pdf/scf14.pdf>
- Feigenbaum, Edward A. 1961. "Soviet cybernetics and computer sciences, 1960." *Communications of the ACM* 4 (12): 566-579.

- Feigenbaum, Edward A. 1968. "Artificial intelligence: themes in the second decade." AI-MEMO-67. Stanford University Department of Computer Science. <https://stacks.stanford.edu/file/druid:cc421zn5789/cc421zn5789.pdf>
- Feigenbaum, Edward A., and Pamela McCorduck. 1984. *The Fifth Generation*. London: Pan Books.
- Feigenbaum, Edward, Pamela McCorduck and Penny Nii. 1988. *The Rise of the Expert Company*. Macmillan.
- Feigenbaum, Edward A., and Julian Feldman, eds. 1963. *Computers and Thought*. New York: McGraw-Hill.
- Ferrucci, David A. 2012. "Introduction to 'This Is Watson'." *IBM Journal of Research and Development* 56 (3.4): 1-1.
- Figure Eight. 2018. "Data Scientist Report 2018." Accessed 13 September 2018. <https://visit.figure-eight.com/rs/416-ZBE-142/images/Data-Scientist-Report.pdf>
- Finn, Chelsea. 2017. "Learning to Learn." *Berkeley Artificial Intelligence Research Blog*. July 18, 2017. Accessed 13 September 2018. <https://bair.berkeley.edu/blog/2017/07/18/learning-to-learn/>
- Fisher, Mark. 1999. "Flatline constructs: gothic materialism and cybernetic theory-fiction." PhD diss., University of Warwick.
- Forbes. 2018. "The 100 largest companies in the world by market value in 2018 (in billion U.S. dollars)." <https://www.statista.com/statistics/263264/top-companies-in-the-world-by-market-value/>
- Forbes, Jennifer. 2019. "'You're good for a girl' – Everyday sexism in tech." *CW Jobs*. August 5, 2019. <https://www.cwjobs.co.uk/blog/everyday-sexism/>
- Ford, Martin. 2009. *The lights in the tunnel: Automation, accelerating technology and the economy of the future*. Acculant Publishing.
- Ford, Martin. 2015. *Rise of the Robots: Technology and the Threat of a Jobless Future*. Basic Books.
- Forsythe, Diana. 2001. *Studying those who study us: An anthropologist in the world of artificial intelligence*. Stanford University Press.
- Frank, Malcolm, Paul Roehrig, and Ben Pring. 2017. *What to do when machines do everything: How to get ahead in a world of AI, algorithms, bots, and Big Data*. John Wiley & Sons.
- Frankish, Keith and William M. Ramsey. 2014. *The Cambridge Handbook of Artificial Intelligence*. Cambridge: Cambridge University Press.
- Frase, Peter. 2015. "Ours to Master." *Jacobin*. <https://jacobinmag.com/2015/03/automation-frase-robots>
- Freund, Karl. 2017. "Google Cloud TPU: Strategic Implications For Google, NVIDIA And The Machine Learning Industry." *Forbes*. May 22, 2017. <https://www.forbes.com/sites/moorinsights/2017/05/22/google-cloud-tpu->

strategic-implications-for-google-nvidia-and-the-machine-learning-industry/#3f2ec8af3af7

- Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The future of employment: how susceptible are jobs to computerisation?" *Technological Forecasting and Social Change* 114: 254-280.
- Freyre, Gilberto. 1966. "Time, Leisure and the Arts: Reflections of a Latin-American On Automation." *Diogenes* 14 (54): 104-115.
- Freud, Sigmund. 1948. *Beyond the Pleasure Principle*. London: Hogarth Press.
- Fuchs, Christian. 2014. *Digital Labour and Karl Marx*. Routledge.
- Fuchs, Christian. 2017. "Reflections on Michael Hardt and Antonio Negri's Book 'Assembly'." *tripleC* 15 (2). <https://www.triple-c.at/index.php/tripleC/article/view/931>
- Fuchs, Christian, and Vincent Mosco. 2015. *Marx in the Age of Digital Capitalism*. Brill.
- Fuchs, Christian, and Nick Dyer-Witheford. 2013. "Karl Marx@ internet studies." *New Media & Society* 15 (5): 782-796.
- Fukuyama, Francis. 2004. "Transhumanism." *Foreign Policy* 144: 42-43.
- Future of Life Institute. n.d. "AI Safety Research." <https://futureoflife.org/ai-safety-research/>
- Gallagher, Sean. 2017. "Expensify sent images with personal data to Mechanical Turkers, calls it a feature." *Ars Technica*. November 27, 2017. <https://arstechnica.com/information-technology/2017/11/expensify-acknowledges-potential-privacy-problem-by-calling-it-a-feature/>
- Gao, Jack. 2017. "China's Wage Growth: How Fast Is the Gain and What Does It Mean?" *Institute for New Economic Thinking*. February 28, 2017. <https://www.ineteconomics.org/perspectives/blog/chinas-wage-growth-how-fast-is-the-gain-and-what-does-it-mean>
- Garcia, Denise. 2014. "The Case against Killer Robots." *Foreign Affairs*.
- Gartner. 2018. "Hype Cycle for Emerging Technologies, 2018." August. [https://blogs.gartner.com/smarterwithgartner/files/2018/08/PR\\_490866\\_5\\_Trends\\_in\\_the\\_Emerging\\_Tech\\_Hype\\_Cycle\\_2018\\_Hype\\_Cycle.png](https://blogs.gartner.com/smarterwithgartner/files/2018/08/PR_490866_5_Trends_in_the_Emerging_Tech_Hype_Cycle_2018_Hype_Cycle.png)
- GE. n.d.a. "Technology Domain: Artificial Intelligence." Accessed May 28, 2018. <https://www.ge.com/research/technology-domains/artificial-intelligence>
- GE. n.d.b. "iFIX." Accessed May 28, 2018. <https://www.ge.com/digital/applications/hmi-scada/ifix>
- Geist, Edward and Andrew J. Lohn, 2018. "How Might Artificial Intelligence Affect the Risk of Nuclear War?" Santa Monica, CA: RAND Corporation. <https://www.rand.org/pubs/perspectives/PE296.html>.
- Geraci, Robert M. 2012. *Apocalyptic AI: Visions of heaven in robotics, artificial intelligence, and virtual reality*. Oxford University Press.

- Gerovitch, Slava. 2004. *From newspeak to cyberspeak: a history of Soviet cybernetics*. MIT Press.
- Gerovitch, Slava. 2011. "Artificial Intelligence with a National Face: American and Soviet Cultural Metaphors for Thought." In *The Search for a Theory of Cognition: Early Mechanisms and New Ideas*, edited by Stefano Franchi and Francesco Bianchini. Rodopi.
- Gershgorn, Dave. 2015. "How Google Aims To Dominate Artificial Intelligence." *Popular Science*. November 9, 2015. <https://www.popsci.com/google-ai>
- Gibbs, Samuel. 2014. "Elon Musk: artificial intelligence is our biggest existential threat." *The Guardian*. October 24, 2014. <https://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat>
- Gillespie, Tarleton. 2010. "The politics of 'platforms'." *New Media & Society* 12 (3): 347-364.
- Goertzel, Ben and Cassio Pennachin, eds. 2007. *Artificial General Intelligence*. New York: Springer.
- Goldfarb, Brent D., Michael D. Pfarrer, and David Kirsch. 2005. "Searching for ghosts: business survival, unmeasured entrepreneurial activity and private equity investment in the dot-com era." *Robert H. Smith School Research Paper* No. 06-027.
- Golumbia, David. 2019. "The Great White Robot God: Artificial General Intelligence and White Supremacy." *Uncomputing* 21. January 2019. [www.uncomputing.org/?p=2077](http://www.uncomputing.org/?p=2077)
- Good, Irving John. 1966. "Speculations concerning the first ultraintelligent machine." In *Advances in Computers* 6: 31-88. Elsevier. [https://www.stat.vt.edu/content/dam/stat\\_vt\\_edu/research/Technical\\_Reports/TechReport05-3.pdf](https://www.stat.vt.edu/content/dam/stat_vt_edu/research/Technical_Reports/TechReport05-3.pdf)
- Goodall, Noah J. 2014. "Machine ethics and automated vehicles." In *Road vehicle automation*, edited by Gereon Meyer and Sven Beiker, 93-102. Springer.
- Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. "Generative adversarial nets." In *Advances in neural information processing systems* 27, edited by Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence and K.Q. Weinberger. 2672-2680.
- Google. n.d. "reCAPTCHA v3: The new way to stop bots." Accessed 7 November 2018. <https://www.google.com/recaptcha/intro/v3beta.html>
- Google. n.d.a. "AI and machine learning products." Accessed May 28, 2018. <https://cloud.google.com/products/ai/>

- Google. n.d.b. "AI solutions." Accessed May 28, 2018. <https://cloud.google.com/solutions/ai/>
- Göranzon, Bo, and Ingela Josefson, eds. 1988. *Knowledge, skill and artificial intelligence*. Berlin: Springer-Verlag.
- Graves, Alex and Kelly Clancy. 2019. "Unsupervised learning: the curious pupil." *DeepMind Blog*. April 10. <https://deepmind.com/blog/unsupervised-learning/>
- Groopman, Jessica. 2017. "Artificial Intelligence: 10 key themes across use cases." *Tractica*. Accessed 6 July 2018. <https://www.tractica.com/resources/white-papers/artificial-intelligence-10-key-themes-across-use-cases/>
- Graham, Stephen D. N. 2005. "Software-sorted Geographies." *Progress in Human Geography* 29 (5): 562-580. <http://www.dourish.com/classes/readings/Graham-SoftwareSortedGeographies-PHG.pdf>
- Graham, Mark, Isis Hjorth, and Vili Lehdonvirta. 2017. "Digital labour and development: impacts of global digital labour platforms and the gig economy on worker livelihoods." *Transfer: European Review of Labour and Research* 23 (2): 135-162.
- Gramsci, Antonio. 1992 [1971]. *Selections from the Prison Notebooks of Antonio Gramsci*, edited and translated by Quintin Hoare and Geoffrey Nowell Smith. New York: International Publishers.
- Granka, Laura A. 2010. "The Politics of Search: A Decade Retrospective." *The Information Society* 26 (5): 364-74. <http://www.australianscience.com.au/research/google/36914.pdf>
- Gray, Mary L., and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Eamon Dolan Books.
- Greenbaum, J. 1979. *In the Name of Efficiency: Management Theory and Shopfloor Practice in Data-Processing Work*. Philadelphia: Temple University Press.
- Greenbaum, J. 1998. "On twenty-five years with Braverman's 'Labour and Monopoly Capital' (Or, How did control and coordination of labor get into software so quickly)". *Monthly Review* 50 (8). <http://go.galegroup.com/ps/i.do?p=AONE&u=lond95336&id=GALE|A53972889&v=2.1&it=r&sid=summon>
- Grier, David Alan. 2013. *When Computers Were Human*. Princeton University Press.
- Grimshaw, Damian, Jo Cartwright, Mat Johnson, and Jill Rubery. 2017. "Market Exposure and the Labour Process: Contradictory Dynamics in Managing Subcontracted Work." in *Academy of Management Proceedings* 2017 (1). New York: Academy of Management.
- Griziotti, Giorgio. 2018. *Neurocapitalism: Technological Mediation and Vanishing Lines*. Minor Compositions.

- Guidotti, Riccardo, Anna Monreale and Dino Pedreschi. 2019. "The AI Black Box Explanation Problem." *ERCIM News* 116. January. <https://ercim-news.ercim.eu/en116/special/the-ai-black-box-explanation-problem>
- Gunning, David. 2017. "Explainable artificial intelligence (xai)." *Defense Advanced Research Projects Agency (DARPA)*.  
<https://www.darpa.mil/attachments/XAIProgramUpdate.pdf>
- Ha, David, and Jürgen Schmidhuber. 2018. "World models." *arXiv preprint arXiv:1803.10122*.
- Hansell, Gregory R. 2011. *H+/-: Transhumanism and its Critics*. Xlibris Corporation.
- Hanson, Robin. 2016. *The Age of Em: Work, Love, and Life when Robots Rule the Earth*. Oxford University Press.
- Hao, Karen. 2019. "Deep learning could reveal why the world works the way it does." *MIT Technology Review*. May 8, 2019. <https://www.technologyreview.com/s/613502/deep-learning-could-reveal-why-the-world-works-the-way-it-does/>
- Haugeland, John. 1989. *Artificial Intelligence: The very idea*. MIT Press.
- Harari, Yuval Noah. 2016. *Homo Deus: A brief history of tomorrow*. Random House.
- Haraway, Donna. 1990. "A manifesto for cyborgs: Science, technology, and socialist feminism in the 1980s." *Feminism/postmodernism*, edited by Linda Nicholson, 190-233. Routledge.
- Hardesty, Larry. 2015. "Automating big-data analysis." *MIT News*. October 16, 2015. <https://news.mit.edu/2015/automating-big-data-analysis-101>
- Hardt, Michael, and Antonio Negri. 2001. *Empire*. Harvard University Press.
- Hardt, Michael, and Antonio Negri. 2005. *Multitude: War and democracy in the age of empire*. Penguin.
- Hardt, Michael, and Antonio Negri. 2009. *Commonwealth*. Harvard University Press.
- Hardt, Michael, and Antonio Negri. 2017. *Assembly*. Oxford University Press.
- Harrington, Stephen, Jeff Moir and J. Scott Allinson. 2018. "The Intelligence Revolution: Future-proofing Canada's workforce." *Deloitte Canada*.  
<https://www2.deloitte.com/ca/en/pages/human-capital/articles/intelligence-revolution.html>
- Harvey, David. 2007. *A brief history of neoliberalism*. Oxford University Press.
- Harvey, David. 2010. *A Companion to Marx's Capital*. Verso.
- Harwell, Drew. 2018. "Google to drop Pentagon AI contract after employee objections to the 'business of war.'" *The Washington Post*. June 1, 2018. [https://www.washingtonpost.com/news/the-switch/wp/2018/06/01/google-to-drop-pentagon-ai-contract-after-employees-called-it-the-business-of-war/?utm\\_term=.f27036684980](https://www.washingtonpost.com/news/the-switch/wp/2018/06/01/google-to-drop-pentagon-ai-contract-after-employees-called-it-the-business-of-war/?utm_term=.f27036684980)



- Hawksworth, John, Richard Berriman and Saloni Goel. 2018. "Will robots really steal our jobs? An international analysis of the potential long term impact of automation." *Pricewaterhousecoopers*. <https://www.pwc.co.uk/economic-services/assets/international-impact-of-automation-feb-2018.pdf>
- Hayles, N. Katherine. 2008. *How we became posthuman: Virtual bodies in cybernetics, literature, and informatics*. University of Chicago Press.
- Hayles, N. Katherine. 2017. *Unthought: The power of the cognitive nonconscious*. University of Chicago Press.
- He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. "Deep residual learning for image recognition." In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770-778. [http://openaccess.thecvf.com/content\\_cvpr\\_2016/papers/He\\_Deep\\_Residual\\_Learning\\_CVPR\\_2016\\_paper.pdf](http://openaccess.thecvf.com/content_cvpr_2016/papers/He_Deep_Residual_Learning_CVPR_2016_paper.pdf)
- Hedberg, Sara Reese. 2002. "DART: revolutionizing logistics planning." *IEEE Intelligent Systems* 17 (3): 81-83.
- Heinrich, Michael. 2012. *An Introduction to the Three Volumes of Karl Marx's Capital*. Aakar Books.
- Heinrich, Michael. 2013. "The 'Fragment on Machines': A Marxian Misconception in the Grundrisse and its Overcoming in Capital." in *In Marx's Laboratory: Critical Interpretations of the Grundrisse*, edited by Riccardo Bellofiore, Guido Starosta and Peter D. Thomas. Leiden/Boston: Brill.
- Highsmith, Jim and Alistair Cockburn. 2001. "Agile Software Development: The Business of Innovation." *Computer* 34 (9): 120-127.
- Hinton, Geoffrey E., J. L. McClelland and D. E. Rumelhart. 1986. *Parallel distributed processing: Explorations in the microstructure of cognition*. MIT Press
- Hinton, Geoffrey, Li Deng, Dong Yu, George Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N. Sainath and Brian Kingsbury. 2012. "Deep neural networks for acoustic modeling in speech recognition." *IEEE Signal processing magazine* 29: 82-97.
- Hitlin, Paul. 2016. "Research in the Crowdsourcing Age, a Case Study." *Pew Research Center*. July 11, 2016. <https://www.pewinternet.org/2016/07/11/research-in-the-crowdsourcing-age-a-case-study/>
- Holloway, John. 2002. *Change the World Without Taking Power: The Meaning of Revolution Today*. Pluto Press.
- Holly, Peter. 2015. "Bill Gates on dangers of artificial intelligence: 'I don't understand why some people are not concerned.'" *The Washington Post*. January 29, 2015. [https://www.washingtonpost.com/news/the-switch/wp/2015/01/28/bill-gates-on-dangers-of-artificial-intelligence-dont-understand-why-some-people-are-not-concerned/?noredirect=on&utm\\_term=.a259c34727b0](https://www.washingtonpost.com/news/the-switch/wp/2015/01/28/bill-gates-on-dangers-of-artificial-intelligence-dont-understand-why-some-people-are-not-concerned/?noredirect=on&utm_term=.a259c34727b0)

- Holzinger, Andreas. 2018. "From Machine Learning to Explainable AI." In *2018 World Symposium on Digital Intelligence for Systems and Machines (DISA)*, 55-66. IEEE.
- Hounshell, David. 2000. "Are Programmers Oppressed by Monopoly Capital, or Shall the Geeks Inherit the Earth? Commentary on Nathan Ensmerger & William Aspray, 'Software as Labor Process'" in *History of Computing: Software Issues*, edited by Ulf Hashagen, Reinhard Keil-Slawik and Arthur L. Norberg, 167-175. New York: Springer-Verlag.
- Hsu, Feng-Hsiung. 2004. *Behind Deep Blue: Building the computer that defeated the world chess champion*. Princeton University Press.
- Huws, Ursula. 2003. *The making of a cybertariat: virtual work in the real world*. Monthly Review Press.
- Huws, Ursula. 2014. *Labor in the global digital economy: The cybertariat comes of age*. NYU Press.
- Hyde, Alan. 2003. *Working in Silicon Valley: Economic and Legal Analysis of a High-Velocity Labor Market*. ME Sharpe.
- Ihara, Kensaku. 2018. "Foxconn plots \$4bn automation push as labor costs bite." *Nikkei Asian Review*. February 24, 2018. <https://asia.nikkei.com/Business/AC/Foxconn-plots-4bn-automation-push-as-labor-costs-bite>
- ImageNet. 2010. "About ImageNet." <http://image-net.org/about-stats>
- IMF. 2017. "Chapter 3: Understanding the Downward Trend in Labor Income Shares." *World Economic Outlook 2017*. <https://www.imf.org/en/Publications/WEO/Issues/2017/04/04/world-economic-outlook-april-2017>
- Indeed.com. 2019a. "Data Scientist Salaries." Accessed July 24, 2019. <https://www.indeed.com/salaries/Data-Scientist-Salaries>
- Indeed.com. 2019b. "Data Engineer Salaries." Accessed July 24, 2019. <https://www.indeed.com/salaries/Data-Engineer-Salaries>
- Indeed.com. 2019c. "Data Analyst Salaries." Accessed July 24, 2019. <https://www.indeed.com/salaries/Data-Analyst-Salaries>
- Information and Communications Technology Council. 2015. "Artificial Intelligence in Canada: Where do we stand?" April 2015. <https://www.ictc-ctic.ca/wp-content/uploads/2015/06/AI-White-paper-final-English1.pdf>
- Intel Brandvoice. 2018. "The Rise in Computing Power: Why Ubiquitous Artificial Intelligence is Now a Reality." *Forbes*. July 17, 2018. <https://www.forbes.com/sites/intelai/2018/07/17/the-rise-in-computing-power-why-ubiquitous-artificial-intelligence-is-now-a-reality/#246769ac1d3f>
- International Data Corporation. 2018. "Worldwide Spending on Cognitive and Artificial Intelligence Systems Will Grow to \$19.1 Billion in 2018, According to New IDC



- Spending Guide.” March 22, 2018.  
<https://www.idc.com/getdoc.jsp?containerId=prUS43662418>
- Jang, Eric. 2017. “What Companies Are Winning the Race for Artificial Intelligence?” *Forbes*. February 24, 2017. <https://www.forbes.com/sites/quora/2017/02/24/what-companies-are-winning-the-race-for-artificial-intelligence/#5ab7c440f5cd>
- Johnson, John M. and Timothy Rowlands. 2012. “The Interpersonal Dynamics of In-Depth Interviewing.” In *The SAGE Handbook of Interview Research*, edited by Jaber F. Gubrium, James A. Holstein, Amir B. Marvasti, and Karyn D. McKinney, 99-113. SAGE.
- Johnson, Matthew, Katja Hofmann, Tim Hutton & David Bignell. 2016. “The Malmo Platform for Artificial Intelligence Experimentation.” In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, 4246-4247. <https://www.ijcai.org/Proceedings/16/Papers/643.pdf>
- Johnston, John. 2008. *The allure of machinic life: cybernetics, artificial life, and the new AI*. MIT Press.
- Jordan, Tim. 2015. *Information Politics: liberation and exploitation in the digital society*. Pluto Press.
- Joy, Bill. 2000. “Why the Future Doesn’t Need Us.” *Wired* 8 (4): 238-262.
- Kania, Elsa. 2017. “AlphaGo and Beyond: The Chinese Military Looks to Future ‘Intelligentized’ Warfare.” *Lawfare Blog*. June 5, 2017. <https://www.lawfareblog.com/alphago-and-beyond-chinese-military-looks-future-intelligentized-warfare>
- Kaplan, Jerrold. 1984. “The industrialization of artificial intelligence: from by-line to bottom line.” *AI Magazine* 5 (2): 51-51.
- Kaplan, Jerry. 2015. *Humans need not apply: A guide to wealth and work in the age of artificial intelligence*. Yale University Press.
- Kaplan, Jerry. 2016. *Artificial Intelligence: What everyone needs to know*. Oxford University Press.
- Karsten, Jack and Darrell M. West. 2015. “Rising minimum wages make automation more cost-effective.” *Brookings Institute Blog*. September 30, 2015. <https://www.brookings.edu/blog/techtank/2015/09/30/rising-minimum-wages-make-automation-more-cost-effective/>
- Kautsky, Karl. 1903. *The Economic Doctrines of Karl Marx*. <https://www.marxists.org/archive/kautsky/1903/economic/index.htm>
- Kaye, Kate. 2019. “These companies claim to provide ‘fair-trade’ data work. Do they?” *MIT Technology Review*. August 7, 2019. <https://www.technologyreview.com/s/614070/cloudfactory-ddd-samasource-imerit-impact-sourcing-companies-for-data-annotation/>
- Kelly, Marlana. 2019. “The Justice Department is working with states on tech investigation, antitrust chief says.” *MSN News*. August 20, 2019.

- <https://www.msn.com/en-us/news/technology/the-justice-department-is-working-with-states-on-tech-investigation-antitrust-chief-says/ar-AAG4Jmg>
- Kelly, Kevin. 2014. "The Three Breakthroughs that have finally unleashed AI on the world." *Wired*. October 27, 2014. <https://www.wired.com/2014/10/future-of-artificial-intelligence/>
- Kisner, James, David Wishnow and Timur Ivannikov. 2017. "IBM: Creating Shareholder Value with AI? Not So Elementary, My Dear Watson." *Jeffries*. July 12, 2017. <https://javatar.bluematrix.com/pdf/fO5xcWjc>
- Kjøsen, Atle Mikkola. 2013. "Do Androids Dream of Surplus Value?" Presented at Mediations 2.5, London, Ontario, January 18, 2013. [www.academia.edu/2455476/Do\\_Androids\\_Dream\\_of\\_Surplus\\_Value](http://www.academia.edu/2455476/Do_Androids_Dream_of_Surplus_Value)
- Klahr, Philip, and Donald A. Waterman. 1986. "Artificial intelligence: A RAND perspective." *AI Magazine* 7 (2): 54-54.
- Kline, Greg. 2003. "Mosaic started Web rush, Internet boom." *The News-Gazette*. April 4, 2003. <http://www.news-gazette.com/news/local/2003-04-20/mosaic-started-web-rush-internet-boom.html>
- Kline, Ronald. 2011. "Cybernetics, automata studies, and the Dartmouth conference on artificial intelligence." *IEEE Annals of the History of Computing* 33 (4): 5-16.
- Knight, Will. 2017. "Reinforcement Learning." *MIT Technology Review*. Accessed November 7, 2018. <https://www.technologyreview.com/s/603501/10-breakthrough-technologies-2017-reinforcement-learning/>
- Knights, David, and Hugh Willmott, eds. 1990. *Labour process theory*. Springer.
- Koch, Christof. 2012. *Consciousness: Confessions of a romantic reductionist*. MIT Press.
- Kockelman, Paul. 2014. "Linguistic anthropology in the age of language automata." *The Cambridge Handbook of Linguistic Anthropology*. Cambridge University Press, 708-733.
- Koehrsen, William. 2018. "Why Automated Feature Learning Will Change the Way you do Machine Learning." *KDNuggets*. <https://www.kdnuggets.com/2018/08/automated-feature-engineering-will-change-machine-learning.html>
- Kolhatkar, Sheelah. 2017. "The Tech Industry's Gender-Discrimination Problem." *The New Yorker*. November 13, 2017. <https://www.newyorker.com/magazine/2017/11/20/the-tech-industrys-gender-discrimination-problem>
- Kolođlugil, Serhat. 2015. "Digitizing Karl Marx: The new political economy of general intellect and immaterial labor." *Rethinking Marxism* 27 (1): 123-137.
- Korsch, Karl. 2013 [1923]. *Marxism and Philosophy*. Verso Trade.
- Kraft, Arnold. 1984. "XCON: An expert configuration system at Digital Equipment Corporation." In *The AI business: The commercial uses of artificial intelligence*, edited by P.H. Winston and K.A. Prenderghast. MIT Press.

- Kraft, P. 1977. *Programmers and Managers: The Routinization of Computer Programming in the United States*. New York: Springer-Verlag.
- Kraft, P. 1979. "The industrialization of computer programming: From programming to 'software production.'" In *Case Studies on the Labor Process*, edited by Andrew S. Zimbalist. New York: Monthly Review Press.
- Kraft, P. 1999. "To control and inspire: US management in the age of computer information systems and global production." In *Rethinking the Labour Process*, edited by Mark Wardell, Thomas L. Steiger and Peter Meiksens. SUNY Press.
- Kraft, Philip, and Steven Dubnoff. 1986. "Job content, fragmentation, and control in computer software work." *Industrial Relations: A Journal of Economy and Society* 25 (2): 184-196.
- Krašovec, Primož. 2018. "The Alien Capital." *Vast Abrupt*. <https://vastabrupt.com/2018/07/11/alien-capital/>
- Krauth, Olivia. 2018. "The 10 tech companies that have invested the most money in AI." Tech Republic. January 12, 2018. <https://www.techrepublic.com/article/the-10-tech-companies-that-have-invested-the-most-money-in-ai/>
- Krishnan, Armin. 2016. *Killer robots: legality and ethicality of autonomous weapons*. Routledge.
- Kuruppu, Dasith. 2019. "Agile development from a programmer's perspective." *Level Up Coding*. January 2, 2019. <https://levelup.gitconnected.com/agile-from-a-developers-perspective-27b23ea665f0>
- Kurz, Heinz D. 2010. "Technical progress, capital accumulation and income distribution in Classical economics: Adam Smith, David Ricardo and Karl Marx." *The European Journal of the History of Economic Thought* 17 (5): 1183-1222.
- Kurzweil, Ray. 1990. *The age of intelligent machines*. Cambridge: MIT Press.
- Kurzweil, Ray. 2000. *The age of spiritual machines: When computers exceed human intelligence*. Penguin.
- Kurzweil, Ray. 2005. *The Singularity is Near: When Humans Transcend Biology*. 2005. Penguin.
- Lacan, Jacques. 1988. *The Seminar of Jacques Lacan*. WW Norton & Company.
- Laitinen, Herbert A. 1970. "Reverberations from the Mansfield Amendment." *Analytical Chemistry* 42 (7): 689-689.
- Land, Nick. 2012. *Fanged Noumena: collected writings 1987-2007*. Urbanomic.
- Land, Nick. 2014. "The Teleological Identity of Capitalism and Artificial Intelligence." Remarks to the participants of the Incredible Machines 2014 Conference, March 8, 2014. Transcription by Jason Adams.
- Land, Nick. 2017. "A quick and dirty guide to accelerationism." *Jacobite*. May 25, 2017. <https://jacobitemag.com/2017/05/25/a-quick-and-dirty-introduction-to-accelerationism>

- Landes, David S. 2003. *The Unbound Prometheus: technological change and industrial development in Western Europe from 1750 to the present*. Cambridge University Press.
- Larson, Jeff, Julia Angwin and Terry Parris Jr.. 2016. "Breaking the Black Box: How Machines learn to be Racist." *ProPublica*. October 19, 2016. <https://www.propublica.org/article/breaking-the-black-box-how-machines-learn-to-be-racist?word=Trump>
- Lazzarato, Maurizio. 1996. "Immaterial Labor." In *Radical Thought in Italy: A potential politics*, edited by Paulo Virno and Michael Hardt, 133-47. University of Minnesota Press.
- Lazzarato, Maurizio. 2014. *Signs and Machines: Capitalism and the production of subjectivity*. Los Angeles: Semiotext(e).
- Leaver, Eric W., and John J. Brown. "Machines without men." *Fortune* 34, no. 11 (1946): 165-204.
- LeCun, Yan, Yoshua Bengio and Geoffrey Hinton. 2015. "Deep Learning." *Nature* 521:436-444.
- Lee, Kai-Fu. 2018. *AI Superpowers: China, Silicon Valley, and the new world order*. Houghton Mifflin Harcourt.
- Lederberg, Joshua. 1987. "How DENDRAL was conceived and born." In *Proceedings of the ACM Symposium on the History of Medical Informatics*. National Library of Medicine. <https://profiles.nlm.nih.gov/ps/access/BBALYP.pdf>
- Lee, Timothy. 2018. "Waymo Jobs: Self-driving Cars Will Destroy a Lot of Jobs – They'll Also Create a Lot." *Ars Technica*, August 24, 2018. <https://arstechnica.com/tech-policy/2018/08/self-driving-cars-will-destroy-a-lot-of-jobs-theyll-alsocreate-a-lot>
- Legg, Shane, and Marcus Hutter. 2007. "A collection of definitions of intelligence." *Frontiers in Artificial Intelligence and applications* 157: 17.
- Le Grand, Rupert. 1948. "Ford Handles with Automation." *American Machinist* 92: 107-109.
- Lenin, Vladimir Ilich. 1920. "Report on the Work of the Council of People's Commissars." December 22, 1920. <http://soviethistory.msu.edu/1921-2/electrification-campaign/communism-is-soviet-power-electrification-of-the-whole-country/>
- Lenin, Vladimir Ilich. 1968. *Collected Works*. Moscow: Progress Publishers.
- Lenin, Vladimir Ilich. 1972 [1914]. *Collected Works*. Progress Publishers, 1972, Moscow, Volume 20, pages 152-154. <https://www.marxists.org/archive/lenin/works/1914/mar/13.htm>
- Leonard, Andrew. 2014. "Tech's toxic political culture: The stealth libertarianism of Silicon Valley bigwigs." *Salon*. June 6,

2014. [https://www.salon.com/2014/06/06/techs\\_toxic\\_political\\_culture\\_the\\_stealth\\_libertarianism\\_of\\_silicon\\_valley\\_bigwigs/](https://www.salon.com/2014/06/06/techs_toxic_political_culture_the_stealth_libertarianism_of_silicon_valley_bigwigs/)
- Levy, Frank, and Richard J. Murnane. 2004. *The new division of labor: How computers are creating the next job market*. Princeton University Press.
- Li, Fei-Fei. 2017. “ImageNet: Where have we been? Where are we going?” *ACM Webinar*. [https://learning.acm.org/binaries/content/assets/learning-center/webinar-slides/2017/imagenet\\_2017\\_acm\\_webinar\\_compressed.pdf](https://learning.acm.org/binaries/content/assets/learning-center/webinar-slides/2017/imagenet_2017_acm_webinar_compressed.pdf)
- Li, Liam and Ameet Talkwalker. 2018. “What is neural architecture search?” *O’Reilly*. December 20, 2018. <https://www.oreilly.com/ideas/what-is-neural-architecture-search>
- Licklider, Joseph CR. 1960. “Man-computer symbiosis.” *IRE transactions on human factors in electronics* 1: 4-11.
- Liebowitz, Jay. 1997. “Worldwide perspectives and trends in expert systems: An analysis based on the three world congresses on expert systems.” *AI Magazine* 18 (2): 115-115.
- Lisboa, P.J.G., and Alfredo Vellido. 2000. “Preface: business applications of neural networks.” In *Business Applications of Neural Networks: The State-of-the-Art of Real-world applications*, edited by P.J.G. Lisboa, Bill Edisbury and Alfredo Vellido. World Scientific.
- Liu, Mingwei, and Chris Smith, eds. 2016. *China at Work: a labour process perspective on the transformation of work and employment in China*. Macmillan International Higher Education.
- Locklin, Scott. 2014. “Neglected machine learning ideas.” *Locklin on Science*. July 22, 2014. <https://scottlocklin.wordpress.com/2014/07/22/neglected-machine-learning-ideas/>
- Lorica, Ben and Mike Loukides. 2018. “Deep automation in machine learning.” *O’Reilly*. December 19, 2018. <https://www.oreilly.com/ideas/deep-automation-in-machine-learning>
- Lu, Yifeng. 2019. “An End-to-End AutoML Solution for Tabular Data at KaggleDays.” *Google AI Blog*. May 9, 2019. <https://ai.googleblog.com/2019/05/an-end-to-end-automl-solution-for.html>
- Lukács, Georg. 1972 [1923]. *History and Class Consciousness: Studies in Marxist Dialectics*. MIT Press.
- lukeprog. 2011. “Better Disagreement.” *LessWrong*. October 24, 2011. <https://www.lesswrong.com/posts/FhH8m5n8qGSSHsAgG/better-disagreement>
- Luo, Zelun, Boya Peng, De-An Huang, Alexandre Alahi, and Li Fei-Fei. 2017. “Unsupervised learning of long-term motion dynamics for videos.” *arXiv preprint arXiv:1701.01821* 2

- Lynch, Shana. 2017. "Andrew Ng: Why AI is the new electricity." *Insights by Stanford Business*. March 11, 2017. <https://www.gsb.stanford.edu/insights/andrew-ng-why-ai-new-electricity>
- Mackenzie, Adrian. 2015. "The production of prediction: What does machine learning want?" *European Journal of Cultural Studies* 18 (4-5): 429-445.
- Mackenzie, Adrian. 2017. *Machine Learners: Archaeology of a Data Practice*. MIT Press.
- Majot, Andrew M., and Roman V. Yampolskiy. 2014. "AI safety engineering through introduction of self-reference into felicific calculus via artificial pain and pleasure." In *Proceedings of the 2014 IEEE International Symposium on Ethics in Science, Technology and Engineering*: 1-6. IEEE.
- Malinovsky, Boris. 2010. *Pioneers of Soviet Computing*. Edited by Anne Fitzpatrick. Translated by Emmanuel Aronie. Published electronically. [http://www.sigcis.org/files/SIGCISMC2010\\_001.pdf](http://www.sigcis.org/files/SIGCISMC2010_001.pdf)
- Manyika, James, Michael Chui, Mehdi Miremadi, Jacques Bughin, Katy George, Paul Willmott and Martin Dewhurst. 2017. "A Future that Works: AI, Automation, Employment, and Productivity." *McKinsey Global Institute*. <https://www.mckinsey.com/~media/mckinsey/featured%20insights/Digital%20Disruption/Harnessing%20automation%20for%20a%20future%20that%20works/MGI-A-future-that-works-Executive-summary.ashx>
- Manzerolle, Vincent R., and Atle Mikkola Kjösen. 2012. "The communication of capital: Digital media and the logic of acceleration." *tripleC: Communication, Capitalism & Critique* 10 (2): 214-229.
- Marcuse, Herbert. (2013 [1964]). *One-dimensional man: Studies in the ideology of advanced industrial society*. Routledge.
- Marazzi, Christian. 2008. *Capital and Language: From the New Economy to the War Economy*, translated by Gregory Conti. Los Angeles: Semiotext [e].
- Marazzi, Christian. 2011. *Capital and Affects: The politics of the language economy*, translated by Giuseppina Mecchia. MIT Press Books.
- Marx, Karl. 1973. *Contribution to a Critique of Political Economy*. Moscow: Progress Publishers.
- Marx, Karl. 1990. *Capital Volume 1*. Penguin.
- Marx, Karl. 1992. *Capital Volume 2*. Penguin.
- Marx, Karl. 1991. *Capital Volume 3*. Penguin.
- Marx, Karl. 1993. *Grundrisse*. Penguin.
- Mason, Paul. 2016. *Postcapitalism: A guide to our future*. Macmillan.
- Mayo, Matthew. 2016. "The Case for Machine Learning in Business." *LinkedIn*. October 3, 2016. <https://www.linkedin.com/pulse/case-machine-learning-business-matthew-mayo/>

- McAfee, Andrew, and Erik Brynjolfsson. 2017. *Machine, platform, crowd: Harnessing our digital future*. WW Norton & Company.
- McCarthy, John, Martin Minsky, N. Rochester and Claude Shannon. 1955. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence." Accessed 6 November 2018. <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>
- McCarthy, John, and Patrick J. Hayes. 1981. "Some philosophical problems from the standpoint of artificial intelligence." In *Readings in artificial intelligence*, edited by Bonnie Lynn Weber and Nils J. Nilsson: 431-450. Morgan Kaufman.
- McClure, Sean. 2018. "GUI-fying the Machine Learning Workflow: Towards Rapid Discovery of Viable Pipelines." *Towards Data Science*. June 25, 2018. <https://towardsdatascience.com/gui-fying-the-machine-learning-workflow-towards-rapid-discovery-of-viable-pipelines-cab2552c909f>
- McCorduck, Pamela. 2004. *Machines Who Think*. Natick, MA: A. K. Peters, Ltd.
- McKinlay, Alan, and Chris Smith, eds. 2009. *Creative Labour: Working in the creative industries*. Palgrave Macmillan.
- McWilliams, Andrew. 2016. "Smart Machines: Technologies and Global Markets Report Overview." BCC Research. <https://www.bccresearch.com/market-research/instrumentation-and-sensors/smart-machines-tech-markets-report.html>
- Meiksins, Peter. 1996. "Work, new technology, and capitalism." *Monthly Review* 48 (3): 99-115.
- Metz, Cade. 2016a. "AI is transforming Google Search. The rest of the web is next." *Wired*. February 4, 2016. <https://www.wired.com/2016/02/ai-is-changing-the-technology-behind-google-searches/>
- Metz, Cade. 2016b. "Building AI is hard - So Facebook is building AI that builds AI." *Wired*. May 6, 2016. <https://www.wired.com/2016/05/facebook-trying-create-ai-can-create-ai/>
- Metz, Cade. 2017a. "Tech Giants are Paying Huge Salaries for Scarce A.I. Talent." *The New York Times*. October 22, 2017. <https://www.nytimes.com/2017/10/22/technology/artificial-intelligence-experts-salaries.html>
- Metz, Cade. 2017b. "Building A.I. that can build A.I." *The New York Times*. November 5, 2017. Accessed 13 September 2018. <https://www.nytimes.com/2017/11/05/technology/machine-learning-artificial-intelligence-ai.html>
- Metz, Cade. 2018a. "A.I. Researchers are making more than \$1 million, even at a Nonprofit." *The New York Times*. April 19, 2018. <https://www.nytimes.com/2018/04/19/technology/artificial-intelligence-salaries-openai.html>

- Metz, Cade. 2018b. "Google sells A.I. for building A.I. (Novices welcome)." *The New York Times*. January 17, 2018. Accessed 13 September 2018. <https://www.nytimes.com/2018/01/17/technology/google-sells-ai.html>
- Metz, Cade. 2018c. "Finally the little guys are moving in on competing with the big chip-makers." *Financial Post*. January 29, 2018. <https://business.financialpost.com/entrepreneur/big-bets-on-ai-open-a-new-frontier-for-chip-startups-too>
- Meyer, Mathias. 2014. "Continuous integration and its tools." *IEEE Software* 31 (3): 14-16.
- Microsoft. 2016. "Democratizing AI: For every person and every organization." <https://news.microsoft.com/features/democratizing-ai/>
- Microsoft. 2018. "The Future Computed: Artificial Intelligence and its role in society." [https://news.microsoft.com/uploads/2018/02/The-Future-Computed\\_2.8.18.pdf](https://news.microsoft.com/uploads/2018/02/The-Future-Computed_2.8.18.pdf)
- Microsoft ML Blog Team. 2018. "How to Implement AI-First Business Models at Scale." *Microsoft Machine Learning Blog*. September 21, 2018. Accessed November 7, 2018. <https://blogs.technet.microsoft.com/machinelearning/2018/09/21/ai-adoption-in-industry-new-report-from-mit-sloan-and-bcg/>
- Miller, Paul D. and Svitlana Matviyenko, eds. 2014. *The Imaginary App*. MIT Press.
- Miller, Ron. 2017. "AWS continues to rule the cloud infrastructure market". *TechCrunch*. October 30, 2017. <https://techcrunch.com/2017/10/30/aws-continues-to-rule-the-cloud-infrastructure-market/>
- Milton, Laurie P. 2003. "An identity perspective on the propensity of high-tech talent to unionize." *Journal of Labor Research* 24 (1): 31-53.
- Minsky, Marvin and Seymour Papert. 1969. *Perceptrons: an introduction to computational geometry*. MIT Press.
- Mittelstadt, Brent Daniel, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter, and Luciano Floridi. 2016. "The ethics of algorithms: Mapping the debate." *Big Data & Society* 3 (2).
- Markoff, John. 2012. "How Many Computers to Identify a Cat? 16,000." *The New York Times*. June 25, 2012. <https://www.nytimes.com/2012/06/26/technology/in-a-big-network-of-computers-evidence-of-machine-learning.html?pagewanted=all>
- Markoff, John. 2016. *Machines of Loving Grace: The quest for common ground between humans and robots*. HarperCollins Publishers.
- Moravec, Hans. 1988. *Mind Children: The future of robot and human intelligence*. Harvard University Press.
- Moravec, Hans. 2000. *Robot: mere machine to transcendent mind*. Oxford University Press.
- More, Max. 1993. "Technological self-transformation." *Extropy* 10 (4-2).



- More, Max, and Natasha Vita-More, eds. 2013. *The Transhumanist Reader: Classical and contemporary essays on the science, technology, and philosophy of the human future*. John Wiley & Sons.
- Morris-Suzuki, Tessa. 1984. "Robots and capitalism." *New Left Review* 147: 109-121.
- Mosco, Vincent, and Catherine McKercher. 2008. *The Laboring of Communication: Will knowledge workers of the world unite?* Rowman & Littlefield.
- Mosco, Vincent. 2015. *To the Cloud: Big data in a turbulent world*. Routledge.
- Mosco, Vincent. 2017. *Becoming Digital: Toward a Post-Internet Society*. Emerald Publishing Limited.
- Mosco, Vincent, and Janet Wasko, eds. 1988. *The Political Economy of Information*. University of Wisconsin Press.
- Moulier-Boutang, Yann. 2012. *Cognitive Capitalism*. Polity.
- Mozur, Paul. 2018. "Inside China's Dystopian Dreams: A.I., Shame and Lots of Cameras." *The New York Times*. July 8, 2018. <https://www.nytimes.com/2018/07/08/business/china-surveillance-technology.html>
- MSV, Janikram. 2017. "In The Era of Artificial Intelligence, GPUs Are The New CPUs." *Forbes*. August 7, 2017. <https://www.forbes.com/sites/janakirammsv/2017/08/07/in-the-era-of-artificial-intelligence-gpus-are-the-new-cpus/#384cb2045d16>
- Muehlhauser, Luke. 2013. "What is AGI?" *Machine Intelligence Research Institute*. August 11, 2013. <https://intelligence.org/2013/08/11/what-is-agi/>
- Mugan, Jonathan. n.d. "Practical Methods for Overcoming The Machine Learning Data Bottleneck." *hackr.io*. <https://hackr.io/event/practical-methods-for-overcoming-the-machine-learning-data-bottleneck-austin>
- Müller, Vincent C., ed. 2016. *Risks of Artificial Intelligence*. CRC Press.
- Muro, Mark, Robert Maxim and Jacon Whiton. 2019. "Automation and Artificial Intelligence: How machines are affecting people and places." *Brookings*. January 24, 2019. <https://www.brookings.edu/research/automation-and-artificial-intelligence-how-machines-affect-people-and-places/>
- Murphy, Kate. 2017. "One Day, a Machine will Smell Whether You're Sick." *The New York Times*. May 1, 2017. Accessed 8 November 2018. <https://www.nytimes.com/2017/05/01/health/artificial-nose-scent-disease.html>
- Mycielski, Jan. 1992. "Games with Perfect Information." In *Handbook of Game Theory with Economic Applications* Volume 2, edited by Robert J. Aumann and Sergiu Hart. Elsevier.
- Naimat, Aman. 2016. *The New Artificial Intelligence Market*. O'Reilly Media.

- Nakashima, Ryan. 2018. "AI's dirty little secret: It's powered by people." *Phys.org*. March 5, 2018. <https://phys.org/news/2018-03-ai-dirty-secret-powered-people.html>
- National Research Council. 1999. *Funding a Revolution: Government support for computing research*. National Academies Press.
- National Research Council Canada. 1991. "Artificial intelligence in Canada: a description by members of the ACAI." NRC No. 31789.
- Negarestani, Reza. 2014. "The Labour of the Inhuman." In *#Accelerate: The Accelerationist Reader*, edited by Robin Mackay and Arven Avanesian: 425-466. Urbanomic.
- Negarestani, Reza. 2015a. "Revolution Backwards: Functional Realization and Computational Implementation." In *Alleys of your Mind: Augmented Intelligence and its Traumas*, edited by Matteo Pasquinelli: 139-156. Luneberg: Meson Press.
- Negarestani, Reza. 2015b. "What is Philosophy? Part One; Axioms and Programs." *e-flux*. November, 2015. <https://www.e-flux.com/journal/67/60702/what-is-philosophy-part-one-axioms-and-programs/>
- Negarestani, Reza. 2016. "What is Philosophy? Part Two: Program and Realizabilities." *e-flux*. January, 2016. <https://www.e-flux.com/journal/69/60608/what-is-philosophy-part-two-programs-and-realizabilities/>
- Negarestani, Reza. 2018. *Intelligence and Spirit*. Urbanomic.
- Negri, Antonio. 1979. "Capitalist domination and working-class sabotage." Online at: <https://antonionegriinenglish.files.wordpress.com/2010/08/capitalist-domination-and-working-class-sabotage.pdf>
- Negri, Antonio. 1989. *The Politics of Subversion: A Manifesto for the Twenty-First Century*. Cambridge: Polity.
- Negri, Antonio. 1991. *Marx Beyond Marx: Lessons on the Grundrisse*. Brooklyn: Autonomedia.
- Negri, Antonio. 1996. "Twenty theses on Marx: Interpretation of the class situation today." In *Marxism Beyond Marxism*, edited by Saree Makdisi, Cesare Casarino and Rebecca Karl, 149-180. Hoboken: Taylor & Francis.
- Negri, Antonio. 2014. "Some Reflections on the #accelerate manifesto." In *#accelerate: The Accelerationist Reader*, edited by Robin Mackay and Arven Avanesian. Urbanomic.
- Nevogt, Dave. 2019. "JIRA Project Management: A How-To Guide for Beginners." *HubStaff*. May 16, 2019. <https://blog.hubstaff.com/jira-project-management-guide-beginners/>
- Newell, Allen, and Herbert A. Simon. 1976. "Computer science as empirical inquiry: symbols and search." *Communications of the ACM* 19 (3): 113-126.

- Newsome, Kirsty, Philip Taylor, Jennifer Bair, and Al Rainnie, eds. 2015. *Putting labour in its place: labour process analysis and global value chains*. Palgrave.
- Ng, Andrew. 2017. Twitter Post. May 1, 2017, 11:04 AM. <https://twitter.com/AndrewYNg/status/859106360662806529>
- Nilsson, Nils. 2005. "Human-Level Artificial Intelligence? Be Serious!" *AI Magazine*. 26 (4): 68-76.
- Nilsson, Nils. 2010. *The Quest for Artificial Intelligence*. Cambridge University Press.
- Noble, David. 1979. *America by Design: Science, technology, and the rise of corporate capitalism*. Oxford University Press.
- Noble, David. 1995. *Progress Without People: New Technology, Unemployment and the Message of Resistance*. Toronto: Between the Lines.
- Noble, David. 2011. *Forces of Production: A Social History of Industrial Automation*. Transaction Publishers.
- Noble, Safiya Umoja. 2018. *Algorithms of Oppression: How search engines reinforce racism*. NYU Press.
- Oakley, Brian, and Kenneth Owen. 1989. *Alvey: Britain's strategic computing initiative*. MIT Press.
- Oberhaus, Daniel. 2018. "Military Contracts Are the Destiny of Every Major Technology Company." *Motherboard*. June 1, 2018. [https://www.vice.com/en\\_us/article/pavzk7/military-contracts-are-the-destiny-of-every-major-technology-company](https://www.vice.com/en_us/article/pavzk7/military-contracts-are-the-destiny-of-every-major-technology-company)
- OECD. 2012. "Chapter 3: Labour Losing to Capital: What Explains the Declining Labour Share?" *OECD Employment Outlook 2012*. [https://www.oecd.org/els/emp/EMO%202012%20Eng\\_Chapter%203.pdf](https://www.oecd.org/els/emp/EMO%202012%20Eng_Chapter%203.pdf)
- Oliveira, Marcia. 2019. "3 Reasons Why AutoML Won't Replace Data Scientists Yet." *KDnuggets*. March 2019. <https://www.kdnuggets.com/2019/03/why-automl-wont-replace-data-scientists.html>
- O'Neil, Cathy. 2017. *Weapons of Math Destruction: How big data increases inequality and threatens democracy*. Broadway Books.
- OpenAI. 2018. "OpenAI Five". *OpenAI Blog*. June 25, 2018. <https://openai.com/blog/openai-five>
- OpenAI. 2019. "OpenAI LP". *OpenAI Blog*. March 11, 2019. <https://openai.com/blog/openai-lp/>
- O'Reilly, Tim. 2017. *WTF?: What's the Future and why It's Up to Us*. Random House.
- Orlikowski, W. J. 1988. "The data processing occupation: Professionalization or proletarianization?" *Research in the Sociology of Work* 4: 95–124.

- Panzieri, Raniero. 2017 [1961]. "The Capitalist Use of Machinery: Marx versus the Objectivists." Available online: <https://libcom.org/library/capalist-use-machinery-raniero-panzieri>
- Panzieri, Raniero. 1965. "Socialist uses of workers' inquiry." Translated by Arianna Bove. originally published in *Spontaneita' e organizzazione. Gli anni dei "Quaderni Rossi" 1959-1964*, a collection of Panzieri's writings edited by S. Merli for BFS Edizioni, Pisa 1994. The English translation was published in Transform. Available online: <http://www.generation-online.org/t/panzieri.htm>
- Partnership on AI. n.d. "Home Page". <https://www.partnershiponai.org>
- Partnership on AI. n.d.a. "About Us". <https://www.partnershiponai.org/about/>
- Partnership on AI. n.d.b. "Meet the Partners". <https://www.partnershiponai.org/partners/>
- Passieri, James. 2015. "GE Expects Predix Software to Do for Factories What Apple's iOS Did for Cell Phones." *The Street*. June 3, 2015. <https://www.thestreet.com/story/13174112/1/ge-expects-predix-software-to-do-for-factories-what-apples-ios-did-for-cell-phones.html>
- Pasquale, Frank. 2016. *The Black Box Society: The secret algorithms that control money and information*. Harvard University Press.
- Pasquinelli, Matteo. 2009. "Google's PageRank algorithm: A diagram of cognitive capitalism and the rentier of the common intellect." In *Deep Search: The politics of search beyond Google*, edited by Konrad Becker and Felix Stalder, 152-163. StudienVerlag.
- Pasquinelli, Matteo. 2015a. "Italian operaismo and the information machine." *Theory, Culture & Society* 32 (3): 49-68.
- Pasquinelli, Matteo, ed. 2015b. *Alleys of your Mind: augmented intelligence and its traumas*. Meson Press.
- Pasquinelli, Matteo. 2016. "Abnormal Encephalization in the Age of Machine Learning." *e-flux* 75. <http://www.e-flux.com/journal/75/67133/abnormal-encephalization-in-the-age-of-machine-learning/>
- Pasquinelli, Matteo. 2017a. "The Automaton of the Anthropocene: On Carbosilicon Machines and Cyberfossil Capital." *South Atlantic Quarterly* 116 (2): 311-326.
- Pasquinelli, Matteo. 2017b. "Machines that morph logic: neural networks and the distorted automation of intelligence as statistical inference." *Glass Bead Journal* 1 (1).
- Pasquinelli, Matteo. (2019 forthcoming). *The Eye of the Master: Capital as Computation and Cognition*. Verso.
- Patel, Siddharth. 2017. "Tech Workers: Friend or Foes?" *Jacobin*. August 25, 2017. <https://www.jacobinmag.com/2017/08/silicon-valley-gentrification-tech-sharing-economy>

- Patki, Neha, Roy Wedge, and Kalyan Veeramachaneni. 2016. "The synthetic data vault." In *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. 399-410. IEEE.
- Patrizio, Andy. 2018. "Top 25 Artificial Intelligence Companies." *Datamation*. April 10, 2018. [www.datamation.com/applications/top-25-artificial-intelligencecompanies.html](http://www.datamation.com/applications/top-25-artificial-intelligencecompanies.html)
- Paul, Sayak. 2018. "Turning Machine Learning Models into APIs in Python". *DataCamp*. October 25, 2018. <https://www.datacamp.com/community/tutorials/machine-learning-models-api-python>
- Pein, Corey. 2012. "Mouthbreathing Machiavellis Dream of a Silicon Reich". *The Baffler*. May 19, 2012. <https://thebaffler.com/latest/mouthbreathing-machiavellis>
- Penrose, Roger. 1989. *The Emperor's New Mind: Concerning computers, minds, and the laws of physics*. Oxford University Press.
- Perez, Carlota. 2009. "The double bubble at the turn of the century: technological roots and structural implications." *Cambridge Journal of Economics* 33 (4): 779-805.
- Perry, Tekla. 2018. "Startup Lanetix Pays US \$775,000 to Software Engineers Fired for Union Organizing." *IEEE Spectrum*. November 12, 2018. <https://spectrum.ieee.org/view-from-the-valley/at-work/tech-careers/startup-lanetix-pays-775000-to-software-engineers-fired-for-union-organizing>
- Peters, Benjamin. 2016. *How Not to Network a Nation: The uneasy history of the Soviet Internet*. MIT Press.
- Peters, Michael A. 2014. "The university in the epoch of digital reason: Fast knowledge in the circuits of cybernetic capitalism." In *Universities in the Flux of Time: An exploration of time and temporality in university life*, edited by Paul Gibbs, Oili-Helena Ylijoki, Carolina Guzmán-Valenzuela and Ronald Barnett, 23-45. Routledge.
- Pham, Thuy T. 2017. "50 Companies leading the AI revolution, detailed." *KDNuggets*. March, 2017. <https://www.kdnuggets.com/2017/03/50-companies-leading-ai-revolution-detailed.html>
- Pichai, Sundar. 2015. "TensorFlow: smarter machine learning, for everyone." *Google Blog*. November 5, 2015. <https://googleblog.blogspot.com/2015/11/tensorflow-smarter-machine-learning-for.html>
- Pitts, Frederick Harry. 2017. "Beyond the Fragment: postoperaismo, postcapitalism and Marx's 'Notes on machines', 45 years on." *Economy and Society* 46 (3-4): 324-345.
- Pitts, Frederick Harry. 2018a. "A crisis of measurability? Critiquing post-operaismo on labour, value and the basic income." *Capital & Class* 42 (1): 3-21.

- Pitts, Frederick Harry. 2018b. *Critiquing capitalism today: New ways to read Marx*. Springer.
- Poster, Mark, and David Savat. 2009. *Deleuze and New Technology*. Edinburgh University Press.
- Predix. n.d. predix.io
- Press, Gil. 2017. "Equifax and SAS Leverage AI And Deep Learning To Improve Consumer Access To Credit." *Forbes*. February 20, 2017. <https://www.forbes.com/sites/gilpress/2017/02/20/equifax-and-sas-leverage-ai-and-deep-learning-to-improve-consumer-access-to-credit/#584aee835a39>
- Purdy, Mark and Paul Daugherty. 2016. "Why Artificial Intelligence is the Future of Growth." *Accenture*. <https://www.accenture.com/us-en/insight-artificial-intelligence-future-growth>
- Ramtin, Ramin. 1991. *Capitalism and automation: Revolution in technology and capitalist breakdown*. Pluto Press.
- RAND Corporation. "A Brief History of RAND." <https://www.rand.org/about/history/a-brief-history-of-rand.html>
- Ranisch, Robert, and Stefan Lorenz Sorgner, eds. 2014. *Post-and transhumanism: An introduction*. Peter Lang Edition.
- Rapp, Nicolas and Brian O'Keefe. 2018. "These 100 Companies Are Leading the Way in A.I." *Fortune*. January 8, 2018. <http://fortune.com/2018/01/08/artificial-intelligence-ai-companies-invest-startups/>
- Raschka, Sebastian. 2016. "How to Explain Machine Learning to a Software Engineer." *KDnuggets*. May 2016. <https://www.kdnuggets.com/2016/05/explain-machine-learning-software-engineer.html>
- Rasmussen, Bente and Birgitte Johansen. 2005. "Trick or treat? Autonomy as Control in Knowledge Work." In *Management, Labour Process and Software Development: Reality Bytes*, edited by Rowena Barret. Routledge.
- Reed, Chris. 2018. "How should we regulate artificial intelligence?" *Philosophical Transactions of the Royal Society: Mathematical, Physical and Engineering Sciences* 376 (2128).
- Ricardo, David. 2001 [1821]. *On The Principles of Political Economy and Taxation*. Kitchener: Batoche Books.
- Ricardo, David. 1951–1973. *The Works and Correspondence of David Ricardo*, 11 vols, edited by Piero Sraffa with collaboration of Maurice H. Dobb. Cambridge: Cambridge University Press.
- Rich, Elaine. 1983. *Artificial Intelligence*. McGraw-Hill.
- Rifkin, Jeremy. 2011. *The third industrial revolution: how lateral power is transforming energy, the economy, and the world*. Macmillan.



- Ritchie, Jane, Jane Lewis and Gillian Elam. 2003. "Designing and Selecting Samples." In *Qualitative Research Practice: A Guide For Social Science Students and Researchers*, edited by Jane Ritchie and Jane Lewis. SAGE.
- Roberts, Sarah T. 2014. "Behind the screen: The hidden digital labor of commercial content moderation." PhD diss., University of Illinois at Urbana-Champaign.
- Robins, Kevin, and Frank Webster. 1988. "Cybernetic capitalism: Information, technology, everyday life." In *The Political Economy of Information*, edited by Vincent Mosco and Janet Wasko, 44-75.
- Roden, David. 2014. *Posthuman life: Philosophy at the edge of the human*. Routledge.
- Roland, Alex, and Philip Shiman. 2002. *Strategic Computing: DARPA and the quest for machine intelligence, 1983-1993*. MIT Press.
- Roose, Kevin. 2013. "Silicon Valley's Anti-Unionism, Now With a Side of Class Warfare." *New York Magazine Intelligencer*. July 11, 2013. <http://nymag.com/intelligencer/2013/07/silicon-valleys-anti-unionism-and-class-warfare.html>
- Ross, Alec. 2016. *The Industries of the Future*. Simon and Schuster.
- Roth, Marcus. 2019. "Artificial Intelligence at the Top 5 US Defense Contractors." *Emerj*. January 3, 2019. <https://emerj.com/ai-sector-overviews/artificial-intelligence-at-the-top-5-us-defense-contractors/>
- Royce, Winston W. 1970. "Managing the development of large software systems: concepts and techniques." In *Proceedings of the 9th International Conference on Software Engineering*, 328-338. IEEE Computer Society Press.
- RT. 2017. "'Whoever leads in AI will rule the world': Putin to Russian children on Knowledge Day." *RT*. September 1, 2017. <https://www.rt.com/news/401731-ai-rule-world-putin/>
- Rubin, Isaak Illich. 1972. *Essays on Marx's theory of value*. Detroit: Black & Red.
- Runfeldt, Melissa. 2017. "Catching up: The Near-Future of Commercial AI." *Insight Data Science*. October 5, 2017. <https://blog.insightdatascience.com/catching-up-the-near-future-of-commercial-ai-95e6fa95069f>
- Russell Stuart, J., and Peter Norvig. 2009. *Artificial Intelligence: A modern approach*. Prentice Hall.
- Sadowski, Jathan. 2018. "Potemkin AI." *Real Life Magazine*. August 6, 2018. <https://reallifemag.com/potemkin-ai/>
- Sammet, Jean E. 1969. *Programming Languages; History and Fundamentals*. Englewood cliffs, N.J.: Prentice Hall.
- Samuel, Arthur L. 1959. "Some studies in machine learning using the game of checkers." *IBM Journal of Research and Development* 3 (3): 210-229.
- Schiller, Dan. 1999. *Digital Capitalism: Networking the global market system*. MIT Press.

- Schiller, Dan. 2014. *Digital Depression: Information technology and economic crisis*. University of Illinois Press.
- Schmidt, Eric, and Jonathan Rosenberg. 2014. *How Google Works*. Hachette UK.
- Schenker, Jason. 2017. *Jobs for Robots*. Prestige Professional Publishing.
- Schwab, Klaus. 2017. *The Fourth Industrial Revolution*. Crown Business.
- Scoville, James G. 2001. "The Taylorization of Vladimir Ilich Lenin." *Industrial Relations: A Journal of Economy and Society* 40 (4): 620-626.
- Science Council Canada (SCC). 1983. "A Workshop on Artificial Intelligence."
- Scharre, Paul. 2018. *Army of None: Autonomous weapons and the future of war*. WW Norton & Company.
- Schomer, Audrey. 2019. "Google ad revenue growth is slowing as Amazon continues eating into its share." *Business Insider*. May 1, 2019. <https://www.businessinsider.com/google-ad-revenue-growth-slows-amazon-taking-share-2019-5>
- Schwaber, Ken, and Mike Beedle. 2002. *Agile software development with Scrum*. Upper Saddle River: Prentice Hall.
- Scrum Alliance. 2018. "State of Scrum 2017-2018." [https://www.scrumalliance.org/ScrumRedesignDEVSite/media/ScrumAllianceMedia/Files%20and%20PDFs/State%20of%20Scrum/2017-SoSR-Final-Version-\(Pages\).pdf](https://www.scrumalliance.org/ScrumRedesignDEVSite/media/ScrumAllianceMedia/Files%20and%20PDFs/State%20of%20Scrum/2017-SoSR-Final-Version-(Pages).pdf)
- Shanahan, Murray. 2015. *The Technological Singularity*. MIT Press.
- Shehabi, Arman, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, Nathaniel Horner, Inês Azevedo, and William Lintner. 2016. "United states data center energy usage report." *Lawrence Berkeley National Laboratory*. <https://eta.lbl.gov/publications/united-states-data-center-energy>
- Sheng, Ellen. 2018. "The shocking truth: In Silicon Valley wages are down for everyone but the top 10%." *Nightly Business Report*. December 3, 2018. <http://nbr.com/2018/12/03/the-shocking-truth-in-silicon-valley-wages-are-down-for-everyone-but-the-top-10/>
- Shodt, Frederik L. 1988. *Inside the Robot Kingdom: Japan, Mechatronics, and the Coming Robotopia*. Kodansha International Ltd.
- Siemens. n.d. "This is Mindsphere!" <https://new.siemens.com/global/en/products/software/mindsphere.html>
- Singer, Daniel. 2019. "This AI Startup Accelerates Machine Learning Model Creation." *StartupHub*. March 21, 2019. <https://www.startuphub.ai/this-ai-startup-accelerates-machine-learning-model-creation/>
- Silver, David, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche,



- Thore Graepel, Demis Hassabis. 2017. "Mastering the game of go without human knowledge." *Nature* 550 (7676): 354.
- Simon, Herbert A., and Allen Newell. 1958. "Heuristic problem solving: The next advance in operations research." *Operations research* 6 (1): 1-10.
- Simon, Herbert A. 1996. *The Sciences of the Artificial*. MIT Press.
- Simondon, Gilbert. 2011. "On the mode of existence of technical objects." *Deleuze Studies* 5 (3): 407-424.
- Simonite, Tom. 2016. "The Missing Link of Artificial Intelligence." *MIT Technology Review*. February 16, 2016. <https://www.technologyreview.com/s/600819/the-missing-link-of-artificial-intelligence/>
- Simonite, Tom. 2017a. "AI software learns to make AI software". *MIT Technology Review*. January 18, 2017. <https://www.technologyreview.com/s/603381/ai-software-learns-to-make-ai-software/>
- Simonite, Tom. 2017b. "Why Google's CEO is excited about automating artificial intelligence". *MIT Technology Review*. May 17, 2017. <https://www.technologyreview.com/s/607894/why-googles-ceo-is-excited-about-automating-artificial-intelligence/>
- Simonite, Tom. 2017c. "For superpowers, artificial intelligence fuels new global arms race." *Wired*. September 8, 2017. <https://www.wired.com/story/for-superpowers-artificial-intelligence-fuels-new-global-arms-race/>
- Simonite, Tom. 2018. "Google's AI Guru Wants Computers to Think More Like Brains." *Wired*. December 12, 2018. <https://www.wired.com/story/googles-ai-guru-computers-think-more-like-brains/>
- Smith, Adam. 1991 [1776]. *Wealth of Nations*. Prometheus Books.
- Smith, Tony. 2000. *Technology and capital in the age of lean production: a Marxian critique of the 'New Economy'*. SUNY Press.
- Smith, Tony. 2009. "The chapters on machinery in the 1861-63 manuscripts." In *Re-reading Marx: New perspectives after the critical edition*, 112-127.
- Smith, Tony. 2013. "The 'general intellect' in the Grundrisse and beyond." *Historical Materialism* 21 (4): 235-255.
- Smith, Reid G., and Robert L. Young. 1984. "The design of the Dipmeter Advisor system." In *Proceedings of the 1984 annual conference of the ACM on the fifth generation challenge*, 15-23. ACM.
- Smith, Paige, Rebecca Greenfield and Jeff Green. 2019. "Gender Pay Reporting May Start in Weeks Across Corporate America." *Bloomberg*. March 6, 2019. <https://www.bloomberg.com/news/articles/2019-03-06/u-s-companies-told-to-report-gender-racial-pay-data-to-eeoc>
- Social Media Collective. 2016. "Critical Algorithm Studies: a Reading List." <https://socialmediacollective.org/reading-lists/critical-algorithm-studies/>

- Sohn-Rethel, Alfred. 1978. *Intellectual and manual labour: A critique of epistemology*. London: Macmillan.
- South, Todd. 2018. "DARPA to use artificial intelligence to help commanders in gray zone conflicts." *Army Times*. March 27, 2018. <https://www.armytimes.com/news/your-army/2018/03/27/darpa-to-use-artificial-intelligence-to-help-commanders-in-gray-zone-conflicts/>
- Spencer, Keith A. and Nicole Karlis. 2019. "Silicon Valley, once a bastion of libertarianism, sees a budding socialist movement." *Salon*. April 11, 2019. <https://www.salon.com/2019/04/11/silicon-valley-once-a-bastion-of-libertarianism-sees-a-budding-socialist-movement/>
- SRI International. n.d. "Shakey the Robot." <https://www.sri.com/work/timeline-innovation/timeline.php?timeline=computing-digital#!&innovation=shakey-the-robot>
- Srnicek, Nick. 2017a. *Platform Capitalism*. John Wiley & Sons.
- Srnicek, Nick. 2017b. "The challenges of platform capitalism: Understanding the logic of a new business model." *Juncture* 23 (4): 254-257.
- Srnicek, Nick, and Alex Williams. 2015. *Inventing the Future: Postcapitalism and a world without work*. Verso Books.
- Stackoverflow. 2018. "Developer Survey Results 2018." <https://insights.stackoverflow.com/survey/2018/#overview>
- Standage, Tom. 2016. "The Return of the Machinery Question." *The Economist*. June 25, 2016. <https://www.economist.com/special-report/2016/06/23/the-return-of-the-machinery-question>
- Statt, Nick and James Vincent. 2018. "Google pledges not to develop AI weapons, but says it will still work with the military." *The Verge*. June 7, 2018. <https://www.theverge.com/2018/6/7/17439310/google-ai-ethics-principles-warfare-weapons-military-project-maven>
- Steels, Luc, and Brice Lepape. 1993. "Knowledge engineering in Esprit." *IEEE Expert* 8 (4): 4-10.
- Steiner, Christopher. 2012. *Automate This: How algorithms took over our markets, our jobs, and the world*. Penguin.
- Steinhoff, James. 2017. "Artificial Intelligence and the Reconfiguration Thesis." Presented at 'Marxism and Revolution Now' Marxist Literary Group Summer Institute on Culture and Society, University of California Davis. June 24-28, 2017. [https://www.academia.edu/34923429/Artificial\\_Intelligence\\_and\\_the\\_Reconfiguration\\_Thesis](https://www.academia.edu/34923429/Artificial_Intelligence_and_the_Reconfiguration_Thesis)
- Steinhoff, James. 2019. "Cognition on Tap: Capital's Theory of AI as Utility". *Digital Culture & Society* 4 (2).
- Stiegler, Bernard. 2010. *For a New Critique of Political Economy*. Polity.

- Stiegler, Bernard. 2017. *Automatic Society: the future of work*. John Wiley & Sons.
- Stone, Peter and Rodney Brooks, Erik Brynjolfsson, Ryan Calo, Oren Etzioni, Greg Hager, Julia Hirschberg, Shivaram Kalyan Krishnan, Ece Kamar, Sarit Kraus, Kevin Leyton-Brown, David Parkes, William Press, AnnaLee Saxenian, Julie Shah, Milind Tambe, and Astro Teller. 2016. "Artificial Intelligence and Life in 2030. One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel." Stanford University. September 2016. <http://ai100.stanford.edu/2016-report>.
- Striphas, Ted. 2015. "Algorithmic Culture." *European Journal of Cultural Studies* 18(4-5): 395-412. <http://ecs.sagepub.com/content/18/4-5/395.abstract>
- Sun, Leo. 2018. "4 New Battleground Markets for Baidu, Alibaba and Tencent." *The Motley Fool*. February 22, 2018. <https://www.fool.com/investing/2018/02/22/4-new-battleground-markets-for-baidu-alibaba-and-t.aspx>
- Sun, Ron. 2014. "Connectionism and Neural Networks." In *The Cambridge Handbook of Artificial Intelligence*, edited by Keith Frankish and William M. Ramsey, 108–27. Cambridge: Cambridge University Press.
- Surber, Regina. 2018. "Artificial Intelligence: Autonomous Technology (AT), Lethal Autonomous Weapons Systems (LAWS) and Peace Time Threats." February 21, 2018. Zurich: ICT4Peace Foundation.
- Susskind, Richard E., and Daniel Susskind. 2015. *The future of the professions: How technology will transform the work of human experts*. Oxford University Press.
- Sutton, Richard S., and Andrew G. Barto. 1998. *Reinforcement Learning: An Introduction*. MIT Press.
- Symonds, Peter. 2019. "Signs of rising levels of workers' struggles in China." *World Socialist Web Site*. February 18, 2019. <https://www.wsws.org/en/articles/2019/02/18/chin-f18.html>
- Tarallo, Bernadette Mary. 1987. "The production of information: an examination of the employment relations of software engineers and computer programmers." PhD Diss, University of California, Davis.
- Taylor, Frederick Winslow. 1914. *The Principles of Scientific Management*. Harper.
- Taylor, Phil, and Sian Moore. 2015. "Cabin crew collectivism: Labour process and the roots of mobilization." *Work, Employment and Society* 29 (1): 79-98.
- Terranova, Tiziana. 2004. *Network Culture: Politics for the information age*. Pluto Press.
- Tegmark, Max. 2017. *Life 3.0: Being human in the age of artificial intelligence*. Knopf.
- Theuwissen, Martijn. 2015. "The different data science roles in the industry." *KDnuggets*. November 2015. <https://www.kdnuggets.com/2015/11/different-data-science-roles-industry.html>
- ThinkGradient. 2019. "Automated Machine Learning – An Overview". *Medium*. January 20, 2019. <https://medium.com/thinkgradient/automated-machine-learning-an-overview-5a3595d5c4b5>

- Think Tank Watch. 2018. "Think Tanks Getting Serious About AI." March 29, 2018. <http://www.thinktankwatch.com/2018/03/think-tanks-getting-serious-about-ai.html>
- Thoburn, Nicholas. 2001. "Autonomous Production? On Negri's New Synthesis." *Theory, Culture & Society* 18 (5): 75-96.
- Thompson, Clive. 2019. *Coders: The Making of a New Tribe and the Remaking of the World*. Penguin.
- Thompson, Paul, and Kendra Briken. 2017. "12 Actually Existing Capitalism: Some Digital Delusions." In *The new digital workplace: How new technologies revolutionise work*, edited by Kendra Briken, Shiona Chillias, Martin Krzywdzinski, 241-263.
- Thompson, Paul, and Chris Smith. 2010. *Working Life: renewing labour process analysis*. Palgrave MacMillan.
- Thompson, Paul and Chris Warhurst. 1998. *Workplaces of the Future*. Palgrave MacMillan.
- Thompson, Paul. 2005. "Foundation and Empire: A critique of Hardt and Negri." *Capital & Class* 29 (2): 73-98.
- Toffler Alvin T. 1980. *The Third Wave*. New York: Bantam Books.
- Toscano, Alberto. 2008. "The open secret of real abstraction." *Rethinking Marxism* 20 (2): 273-287.
- Toscano, Alberto. 2009. "Chronicles of insurrection: Tronti, Negri and the subject of antagonism." *Cosmos and History: The Journal of Natural and Social Philosophy* 5 (1): 76-91.
- Toscano, Alberto. 2011. "Logistics and opposition." *Mute* 3 (2): 30-41.
- Toussaint, Eric. 2009. "Adam Smith is closer to Karl Marx than those showering praise on Smith today." *Committee for the Abolition of Illegitimate Debt*. June 12, 2009. <http://www.cadtm.org/Adam-Smith-is-closer-to-Karl-Marx#nh4>.
- Toynbee, Arnold. 1997. *Lectures on the Industrial Revolution in England: Popular addresses, notes and other fragments*. Rivingtons.
- Tractica. 2016. "Executive Summary: Artificial Intelligence Market Forecasts." <https://www.tractica.com/wp-content/uploads/2016/08/MD-AIMF-3Q16-Executive-Summary.pdf>
- Tractica. 2019. "Revenues from the artificial intelligence (AI) software market worldwide from 2018 to 2025 (in billion U.S. dollars)." *Statista*. <https://www.statista.com/statistics/607716/worldwide-artificial-intelligence-market-revenues/>
- TransmogriAI. n.d. "About." <https://transmogri.ai/>
- Trefis Team. 2019. "Can Google Cloud Revenues Reach \$20 Billion by 2020?" *Forbes*. April 29, 2019. <https://www.forbes.com/sites/greatspeculations/2019/04/29/why-cloud-could-be-a-game-changer-for-google/#2b13fef21694>

- Tronti, Mario. 1963. "Social Capital." Available online: <http://zerowork.org/TrontiSocialCapital.htm>
- Tronti, Mario. 1964. "Lenin in England." *Classe Operaia* 1. <https://www.marxists.org/reference/subject/philosophy/works/it/tronti.htm>
- Tronti, Mario. 1966. *Operai e capitale*. Einaudi Editore. <https://libcom.org/library/operai-e-capitale-mario-tronti>
- T.S.. 2016. "Why artificial intelligence is enjoying a renaissance." *The Economist*. July 15, 2016. <https://www.economist.com/the-economist-explains/2016/07/15/why-artificial-intelligence-is-enjoying-a-renaissance?zid=291&ah=906e69ad01d2ee51960100b7fa502595>
- Trump, Donald J. 2019. "Executive Order on Maintaining American Leadership in Artificial Intelligence." *Whitehouse.gov*. February 11, 2019. <https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/>
- Turing, Alan Mathison. 2004. *The Essential Turing: Seminal Writings in Computing, Logic, Philosophy*, edited by Jack Copeland. Oxford University Press.
- Turner, Matt. 2016. "The CEO of United Technologies just let slip an unintended consequence of the Trump-Carrier jobs deal." *Business Insider*. December 5, 2016. <https://www.businessinsider.com/united-tech-ceo-says-trump-deal-will-lead-to-more-automation-fewer-jobs-2016-12>
- Ure, Andrew. 1835. *The Philosophy of Manufactures, or an exposition of the Scientific, Moral, and Commercial Economy of the Factory System of Great Britain*. London, Charles Knight.
- US Bureau of Labor Statistics. 1993. "Labor Composition and U.S. Productivity Growth, 1948-90." [https://www.bls.gov/mfp/labor\\_composition.pdf](https://www.bls.gov/mfp/labor_composition.pdf)
- U.S. Equal Employment Opportunity Commission. 2016. "Diversity in High Tech." <https://www.eeoc.gov/eeoc/statistics/reports/hightech/>
- Van Otterlo, Martjin and Marco Wiering. 2012. "Reinforcement Learning and Markov Decision Processes." In *Reinforcement Learning: State-of-the-Art*, edited by Marco Wiering and Martijn van Otterlo, 3-42. Berlin: Springer-Verlag.
- Vardalas, John N. 2001. *The computer revolution in Canada: building national technological competence*. MIT Press.
- Vellido, Alfredo, Paulo J.G. Lisboa, and J. Vaughan. 1999. "Neural networks in business: a survey of applications (1992-1998)." *Expert Systems with applications* 17 (1): 51-70.
- Virno, Paolo. 2001. "On General Intellect." <https://libcom.org/library/on-general-intellect-paulo-virno>
- Virno, Paolo. 2003. *A Grammar of the Multitude*. Los Angeles: Semiotext(e).

- Vorhies, William. 2017a. "Data Scientists Automated and Unemployed by 2025 – Update!" *Data Science Central*. July 17, 2017. <https://www.datasciencecentral.com/profiles/blogs/data-scientists-automated-and-unemployed-by-2025-update>
- Vorhies, William. 2017b. "More on Fully Automated Machine Learning." *Data Science Central*. August 15, 2017. <https://www.datasciencecentral.com/profiles/blogs/more-on-fully-automated-machine-learning>
- Vorhies, William. 2018. "Next Generation Automated Machine Learning (AML)." *Data Science Central*. April 3, 2018. <https://www.datasciencecentral.com/profiles/blogs/next-generation-automated-machine-learning-aml>
- Wallén, Johanna. 2008. *The History of the Industrial Robot*. Linköping University Electronic Press.
- Wang, Pei. 2008. "What Do You Mean by 'AI'?" *Artificial General Intelligence*, edited by Ben Goertzel and Cassio Pennachin, 362-373. Springer.
- Wardell, Mark, Thomas L. Steiger and Peter Meiksins (eds). 1999. *Rethinking the Labor Process*. New York: State University of New York Press.
- Webster, Frank. 2014. *Theories of the Information Society*. Routledge.
- Wiener, Norbert. 1948. *Cybernetics*. Technology Press.
- Wiener, Norbert. 1949. "Father of Cybernetics Norbert Wiener's Letter to UAW President Walter Reuther." *LibCom*. <https://libcom.org/history/father-cybernetics-norbert-wieners-letter-uaw-president-walter-reuther>
- Williams, Alex, and Nick Srnicek. 2014. "#accelerate: Manifesto for an accelerationist politics." In *#accelerate: the Accelerationist Reader*, edited by Robin Mackay and Arven Avanesian. Urbanomic.
- Womack, James P., James P. Womack, Daniel T. Jones, and Daniel Roos. 1990. *The Machine that changed the world*. Simon and Schuster.
- World Economic Forum. 2016. "The Future of Jobs: Employment, Skills, and Workforce Strategy for the Fourth Industrial Revolution." In *Global Challenge Insight Report*. Geneva: World Economic Forum.
- Wood, Stephen. 1982. *The Degradation of Work?: skill, deskilling, and the labour process*. Hutchinson Radius.
- Woyke, Elizabeth. 2017. "General Electric Builds an AI Workforce." *MIT Technology Review*. June 27, 2017. <https://www.technologyreview.com/s/607962/general-electric-builds-an-ai-workforce/>
- Wright, Steve. 2002a. "The historiography of the mass worker." *The Commoner* 5. <https://libcom.org/library/historiography-mass-worker-steve-wright>
- Wright, Steve. 2002b. *Storming Heaven*. London: Pluto Press.



- Vance, Ashlee. 2014. "The Race to Buy the Human Brains Behind Deep Learning Machines." *Bloomberg*. January 27, 2014. <https://www.bloomberg.com/news/articles/2014-01-27/the-race-to-buy-the-human-brains-behind-deep-learning-machines>
- Vassallo, Trae, Ellen Levy, Michele Madansky, Hillary Mickell, Bennett Porter, Monica Leas and Julie Oberweis. 2015. "Elephant in the Valley." <https://www.elephantinthevalley.com/>
- Vast Abrupt. 2018. "Ideology, Intelligence and Capital: An Interview with Nick Land." August 15, 2018. <https://vastabrupt.com/2018/08/15/ideology-intelligence-and-capital-nick-land/>
- Vercellone, Carlo. 2005. "The hypothesis of cognitive capitalism." No. halshs-00273641. HAL. <https://halshs.archives-ouvertes.fr/halshs-00273641/>
- Vercellone, Carlo. 2007. "From formal subsumption to general intellect: Elements for a Marxist reading of the thesis of cognitive capitalism." *Historical materialism* 15 (1): 13-36.
- Victor, Bret. 2013. "References for 'The Future of Programming'." *Worrydream.com*. <http://worrydream.com/dbx/>
- Vincent, James. 2018. "The world's most valuable AI startup is a Chinese company specializing in real-time surveillance." *The Verge*. April 11, 2018. <https://www.theverge.com/2018/4/11/17223504/ai-startup-sensetime-china-most-valuable-facial-recognition-surveillance>
- Vinge, Vernor. 2013. [1993]. "Technological singularity." In *The Transhumanist Reader: Classical and Contemporary Essays on the Science, Technology, and Philosophy of the Human Future*, edited by Max More and Natasha Vita-More, 365-375. John Wiley & Sons.
- Barker, Virginia, E. and Dennis E. O'Connor. 1989. "Expert systems for configuration at Digital: XCON and beyond." *Communications of the ACM* 32 (3): 298-318.
- Wakabayashi, Daisuke. 2017. "Google Fires Engineer Who Wrote Memo Questioning Women in Tech." *The New York Times*. August 7, 2017. <https://www.nytimes.com/2017/08/07/business/google-women-engineer-fired-memo.html>
- Walker, Jon. 2019. "Machine Learning in Manufacturing – Present and Future Use-Cases." *Emerj*. August 13, 2019. <https://emerj.com/ai-sector-overviews/machine-learning-in-manufacturing/>
- Wallach, Wendell, and Colin Allen. 2008. *Moral Machines: Teaching robots right from wrong*. Oxford University Press.
- Wang, Jane X., Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z. Leibo, Remi Munos, Charles Blundell, Dharsan Kumaran and Matt Botvinick. 2017. "Learning to reinforcement learn." *ArXiv*. <https://arxiv.org/abs/1611.05763>
- Warwick, Kevin. 2013. *Artificial Intelligence: The Basics*. Routledge.

- Wistuba, Martin. 2018. "Using AI to Design Deep Learning Architectures." *IBM Research Blog*. September 4, 2018. <https://www.ibm.com/blogs/research/2018/09/ai-design-deep-learning/>
- Warren, Elizabeth. 2019. "Here's how we can break up Big Tech." *Medium*. March 8, 2019. <https://medium.com/@teamwarren/heres-how-we-can-break-up-big-tech-9ad9e0da324c>
- Wei-Haas, Maya. 2019. "What the new artificial intelligence initiative does – and doesn't – mean." *National Geographic*. February 12, 2019. <https://www.nationalgeographic.com/science/2019/02/what-trump-american-artificial-intelligence-initiative-ai-does-doesnt-mean/>
- West, Sarah Meyes, Meredith Whittaker and Kate Crawford. 2019. "Discriminating Systems: Gender, Race and Power in AI." *AI Now Institute*. April 2019. <https://ainowinstitute.org/discriminatingystems.html>.
- White, Jeffrey, ed. 2015. *Rethinking machine ethics in the age of ubiquitous technology*. IGI Global.
- Winfield, Alan. 2012. *Robotics: A Very Short Introduction*. MIT Press.
- Witheyford, Nick. 1994. "Autonomist Marxism and the information society." *Capital & Class* 18 (1): 85-125.
- Working Partnerships USA. 2016. "Tech's Invisible Workforce." <https://www.wpusa.org/files/reports/TechsInvisibleWorkforce.pdf>
- Wright, E.O. 1981. "The Value Controversy and Social Research." In *The Value Controversy*, edited by Ian Steedman. London: Verso.
- Wu, Jimmy. 2019. "Optimize What?" *Commune*. March 15, 2019. <https://communemag.com/optimize-what/>
- Yao, Mariya. 2019. "A quick overview of enterprise AutoML solutions." *Topbots*. March 31, 2019. <https://www.topbots.com/automl-solutions-overview/>
- Yegulalp, Serdar. 2017. "Google's machine learning cloud pipeline explained." *Infoworld*. May 19, 2017. <https://www.infoworld.com/article/3197405/tpus-googles-machine-learning-pipeline-explained.html>
- Yourdon, Edward. 2004. *Death March*. Prentice Hall Professional Technical Reference.
- Yu, Miaomiao. 2017. "The Humans Behind Artificial Intelligence." *Synced*. April 30. <https://syncedreview.com/2017/04/30/the-humans-behind-artificial-intelligence/>
- Zilis, Shivon and James Cham. 2016. "The current state of machine intelligence 3.0." O'Reilly. November 7, 2016. <https://www.oreilly.com/ideas/the-current-state-of-machine-intelligence-3-0>
- Zimmerman, Andrew. 1997. "The Ideology of the Machine and the Spirit of the Factory: Remarx on Babbage and Ure." *Cultural Critique* 37: 5-29.



- Zittrain, Jonathan. 2019. "The Hidden Costs of Automated Thinking." *The New Yorker*. July 23, 2019. <https://www.newyorker.com/tech/annals-of-technology/the-hidden-costs-of-automated-thinking>
- Žmolek, Michael Andrew. 2013. *Rethinking the Industrial Revolution: five centuries of transition from agrarian to industrial capitalism in England*. Brill.
- Zöllner, Marc-André, and Marco F. Huber. 2019. "Survey on Automated Machine Learning." *arXiv preprint*. arXiv:1904.12054.
- Zoph, Barret and Quoc V. Le. 2017. "Neural Architecture Search with Reinforcement Learning." Under review as conference paper at ICLR 2017. <https://arxiv.org/pdf/1611.01578.pdf>
- Zuboff, Shoshana. 2015. "Big other: surveillance capitalism and the prospects of an information civilization." *Journal of Information Technology* 30 (1): 75-89.

## Appendices

### Appendix 1: Recruitment Blurb

Dear [company/individual],

I'm a PhD candidate in the Faculty of Information and Media Studies at The University of Western Ontario. My doctoral research is concerned with the labour practices and experiences of people working in and around the field of *machine learning* in a commercial setting. I am looking for volunteers who would be willing to participate in an approximately 1-hour interview with me on the topic. I'm interested in speaking with developers, CEOs and anyone else. This research is funded in part by a grant from the Social Sciences and Humanities Research Council of Canada (SSHRC).

Would your company be willing to distribute a call for interviewees to your employees? This could be done via a forwarded email, or by posting a recruitment poster in the workplace (e.g. break room). If this might be possible, or if you'd like more information, please contact me and I will be happy to provide you with the official Letter of Information. I have attached a recruitment poster.

Thank you,

James Steinhoff

PhD Candidate

Faculty of Information and Media Studies

The University of Western Ontario

## Appendix 2: Recruitment Poster

### PARTICIPANTS NEEDED FOR RESEARCH IN THE PRODUCTION OF MACHINE LEARNING TECHNOLOGY

We are looking for volunteers to take part in a study of the production of machine learning technology who are employees or executives of a company which produces machine learning technology

If you are interested and agree to participate you would be asked to take part in an individual interview.

Your participation would involve 1 session which will be about 60 minutes long.

You will not be compensated for your participation.

For more information about this study, or to volunteer for this study, please contact:

Primary Investigator

Nick Dyer-Witthford, PhD.

Faculty of Information and Media Studies

The University of Western Ontario

Research Assistant

James Steinhoff, PhD Candidate

Faculty of Information and Media Studies

## Appendix 3: Letter of Information and Consent

### Letter of Information and Consent: The Production of Machine Learning Technology

Dr. Nick Dyer-Witheford, PhD, Information and Media Studies

The University of Western Ontario

James Steinhoff, PhD Candidate, Information and Media Studies

The University of Western Ontario

#### Invitation to Participate

You are being invited to participate in this research study about the production of machine learning technology because you work in this field.

Why is this study being done?

The goal of this project is to understand how machine learning technologies are produced, focusing on skills required, tools employed, labour processes and organizational systems.

How long will you be in this study?

It is anticipated that the interview will take one hour in one session.

What are the study procedures?

If you agree to participate, you will be asked to engage in a one-on-one interview with a researcher. The interview will be conducted at your office or an agreeable public space (eg. library, café) of your preference. This interview will be audio recorded for later transcription, after which it will be destroyed. The transcriptions will be anonymized. If you do not wish to be audio recorded, the researcher will take notes by hand. There will be a total of 15 participants in the study.

What are the risks and harms of participating in this study?

There is a possible risk that critical opinions of your workplace or co-workers expressed by you could be detrimental to your career if they were disseminated. However, all of your responses will be anonymized and will not be attributable to you.

What are the benefits of participating in this study?

A possible benefit to you may be an improved understanding of labour conditions in machine learning. A possible benefit to society may be a better understanding of an increasingly important sector of high-technology work.

Can participants choose to leave the study?

You may withdraw from the study at any time. If you decide to withdraw from the study, you have the right to request withdrawal of information collected about you. If you wish to have your information removed please let the researcher know.

How will participants' information be kept confidential?

The audio recordings of interviews will be stored on an encrypted file and password-protected hard drive which only the primary investigator and research assistant will have access to.

The recordings will be destroyed after transcription and the transcripts will be anonymized and stored on an encrypted and password-protected hard drive.

If the results of the study are published, your name will not be used.

Representatives of The University of Western Ontario Non-Medical Research Ethics Board may require access to your study-related records to monitor the conduct of the research.

While we do our best to protect your information there is no guarantee that we will be able to do so. If there is data collected during the project which is required by law to report we have a duty to report.

The researcher will keep any personal information about you in a secure and confidential location for a minimum of 5 years. A list linking your study number with your name will be kept by the researcher in a secure place, separate from your study file.

Are participants compensated to be in this study?

You will not be compensated for your participation in this research.

What are the rights of participants?

Your participation in this study is voluntary. You may decide not to be in this study. Even if you consent to participate you have the right to not answer individual questions or to withdraw from the study at any time. If you choose not to participate or to leave the study at any time it will have no effect on your employment status.

We will give you new information that is learned during the study that might affect your decision to stay in the study.

You do not waive any legal right by signing this consent form.

Whom do participants contact for questions?

If you have questions about this research study please contact:

James Steinhoff, PhD Candidate

Dr. Nick Dyer-Witthford, PhD

If you have any questions about your rights as a research participant or the conduct of this study, you may contact:

The Office of Human Research Ethics

Appendix 4 : Ethics Approval



Western University Non-Medical Research Ethics Board  
NMREB Delegated Initial Approval Notice

**Principal Investigator:** Dr. Nicholas Dyer-Witford  
**Department & Institution:** Information and Media Studies/Faculty of Information & Media Studies, Western University

**NMREB File Number:** 109130  
**Study Title:** The Production of Machine Learning Technology

**NMREB Initial Approval Date:** May 01, 2017  
**NMREB Expiry Date:** May 01, 2018

**Documents Approved and/or Received for Information:**

Document Name	Comments	Version Date
Recruitment Items		2017/02/25
Recruitment Items		2017/02/25
Western University Protocol	Received April 17, 2017	
Letter of Information & Consent	Received April 17, 2017	
Instruments	Interview Questions Received April 17, 2017	

The Western University Non-Medical Research Ethics Board (NMREB) has reviewed and approved the above named study, as of the NMREB Initial Approval Date noted above.

NMREB approval for this study remains valid until the NMREB Expiry Date noted above, conditional to timely submission and acceptance of NMREB Continuing Ethics Review.

The Western University NMREB operates in compliance with the Tri-Council Policy Statement Ethical Conduct for Research Involving Humans (TCPS2), the Ontario Personal Health Information Protection Act (PHIPA, 2004), and the applicable laws and regulations of Ontario.

Members of the NMREB who are named as Investigators in research studies do not participate in discussions related to, nor vote on such studies when they are presented to the REB.

The NMREB is registered with the U.S. Department of Health & Human Services under the IRB registration number IRB 00000941.

Ethics Officer \_\_\_\_\_ Chair or delegated board member

EO: Erika Basile \_\_\_ Grace Kelly  Katelyn Harris \_\_\_ Nicola Morphet \_\_\_ Karen Gopaul \_\_\_

## Appendix 5 : Interview Questions

### 1 - Labour process/organization

What is your position?

What is your educational background?

What does your company produce?

Can you describe a typical day at work?

How many hours do you typically work per day? Per week?

What is the organizational structure of your workplace?

Do you tend to work alone or in groups?

What development process [and business process?] does your workplace employ? -agile vs life cycle (waterfall) processes?

How does it work?

How are you involved?

What programs, languages and tools do you use?

What is the place of open-source in your work? Github etc.

What free services do you use?

Have there been any significant changes to your typical activities at work during your employment?

What is your company's long-term goal?

What are the biggest problems you face in your work?



## 2 - ML

How do you and your company use ML?

How do you use other kinds of AI?

Have you worked in any other ML companies prior to this one?

Have you worked in other fields of software production? How is machine learning different?

What are the most exciting prospects/applications for machine learning (at your work and generally)?

What are the biggest problems for machine learning?

What are the most salient risks or dangers presented by machine learning?

Do you think of ML as an automation technology?

Are aspects of your work automated now? Do you think aspects of your work will be automated in the future?

Do you work on personal machine learning projects when not at work?

## 3 - Use of ML to make ML

Do you employ machine learning or artificial intelligence technology in your work? If so, how? If not, how do you think it could or will be?

Are you aware of Google's/Facebook's (or any other) employment of machine learning technology in the production of machine learning technology? What do you think of it?

## 4 - Speculations

What do you think of the Computational Creativity paradigm?

What do you think are the prospects of AGI and ASI?

What is your opinion of the Technological Unemployment thesis? Or does new technology generate new jobs?

## Curriculum Vitae

**Name:** James Steinhoff

**Post-secondary Education and Degrees:** University of Windsor  
Windsor, Ontario, Canada  
2005-2010 B.A. English Literature and Philosophy

University of Windsor  
London, Ontario, Canada  
2011-2011 M.A. Philosophy

The University of Western Ontario  
London, Ontario, Canada  
2014-2019 Ph.D. Media Studies

**Honours and Awards:** Social Science and Humanities Research Council (SSHRC)  
Doctoral Scholarship (\$40,000)  
2017-2019

**Related Work Experience** Teaching Assistant  
The University of Western Ontario  
2014-2019

**Publications:**

Dyer-Withford, Nick, Atle Mikkola Kjösen and James Steinhoff. 2019. *Inhuman Power: Artificial Intelligence and the Future of Capitalism*. London: Pluto Press.

Steinhoff, James. 2019. "Cognition on Tap: Capital's Theory of AI as Utility." *Digital Culture & Society* 4 (2).