

UNIVERSITY OF BERGEN

MASTERS THESIS

**Opportunity and risk are twins.
An investigation into algorithms.**

Author:
Katherine WIDOMSKI

Supervisor:
Dr. Alf Erling RISA

*A thesis submitted in fulfilment of the requirements
for the degree of Master of Economics*

in the

Department of Economics

December 1, 2018

Declaration of Authorship

I, Katherine WIDOMSKI, declare that this thesis titled, “Opportunity and risk are twins.

An investigation into algorithms.” and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Katherine Widomski

Date:

1.12.18



"Unfortunately, a tiny percentage of the drones are opposed to violence."

FIGURE 1: New Yorker Magazine, August 24th 2018.
Artist: Farley Katz.

UNIVERSITY OF BERGEN

Abstract

Department of Economics

Master of Economics

**Opportunity and risk are twins.
An investigation into algorithms.**

by Katherine WIDOMSKI

Algorithms will free up people's time to do other things, much like computers have freed people from paper calculations. The depth and breadth of the application of algorithms is best understood through a task-based automation model. An algorithm cannot replace all of the duties of a certain, perhaps even cognitively taxing job; however, an algorithm can be used to automate many of the so-called individual tasks. Their ability to replace physical tasks is more limited. An algorithm can also be used to perform certain tasks to a higher quality in a shorter period of time than a person. Algorithms may globally become the ultimate productivity boost for advanced economies. Hence, it is recommended that people learn to work with algorithms, instead of against them. In addition, by combining human creativity with algorithmic power, a new era can be ushered into existence. . . .

Acknowledgements

While I was working on my master's thesis, I was also working full time and taking care of my wonderful daughter Vera. Consequently, this thesis would not have been completed without Norway's childcare guarantee and the help of many people.

Firstly, I would like to thank my regular babysitter, Kasia Sawa, for looking after Vera for so many nights of the week, so that I could go out and write.

Next, I would like to thank Sjøsprøyten barnehage, for taking such excellent care of my daughter during the day. The team at Sjøsprøyten barnehage is wonderful. I have become friends with many of the families who attend there, because they welcomed me into their community.

I would like to thank the loving Berle family for taking Vera to gymlek, for feeding us, for looking after us, and for making me feel part of the Bergen community. Your care and love continued to motivate me through many of my difficult days.

Without the world's best duo, Hanne Petrine Grotle and Erik Førde, I do not know how I would have made it through November! Thank you for taking Vera overnight, keeping her safe and keeping her amused while Andreas was away.

I sincerely thank my family back home in Australia, for encouraging reading, critical thinking, knowledge, and the pursuit of something bigger.

I thank my walking jokers, Ana and Mariana. Where would I be without you?

I thank the Steinsland family for taking me under their wing during the most difficult period of my life; I also thank you for always welcoming me into your home and for giving Vera and I a safe place to be.

I thank my fantastic colleagues at Vizrt for their patience while I wrote this!

Outside of my gratitude for the entire village of Bergen, who loved and cared for me and my daughter, I would also like to thank the people who helped me with my intellectual pursuits. Thanks go to my supervisor, for his open-mindedness, curiosity, patience, eagerness to go on wild philosophical journeys, and attempts to bring me back to reality. It is good to have had access to your great mind.

Thank you Brendan Markey-Towler, for chatting to me, for encouraging me, for questioning me, for explaining some of the methodology, for positing interesting notions, and for showing me different sides to arguments.

Thank you Dr. Pamela Kaval (aka Dr. Pam) for your excellent comments and critiques.

Finally, thank you to my husband Andreas for his love. ...

Contents

Declaration of Authorship	iii
Abstract	vii
Acknowledgements	ix
1 Why tasks? Why algorithms? Why work at all?	1
1.1 A short introduction	1
1.2 Working papers	3
2 A model of task based automation	5
2.0.1 On algorithms	5
2.0.2 Displacement?	6
2.1 Restrepo & Aemođlu’s model, with comments	7
2.1.1 Korinek & Stiglitz	11
2.1.2 A note on capital: will algorithms be considered capital, labour, or technology?	13
3 The What and the How:	15
3.1 What we talk about, when we talk about algorithms	15
3.1.1 Artificial Intelligence	15
3.1.2 Machine Learning	16
3.1.3 Deep Learning	19
3.1.4 Reinforcement Learning	19
3.1.5 A comment on the term "learning"	20
3.1.6 Natural Language Processing	20
3.1.7 Computer Vision	21
3.2 What drives these algorithms?	21
3.2.1 Disruption	23
3.2.2 Demand for labour in the world of algorithms	24
3.2.3 Let’s get hyper-specialized	24
3.2.4 What about wages?	24
3.2.5 What about productivity?	25
3.2.6 What about job polarization?	25
3.2.7 Will we tax robots?	25
3.3 Are we headed towards Singularity ?	26

3.4	Let's talk general Empirics	26
3.5	The effects of algorithms on three professions	27
3.5.1	The effect on the medical profession	28
3.5.2	The effect on the financial profession	28
3.5.3	The effect on the legal profession	29
3.5.4	But we are lacking data!	30
4	Policy for the uncanny valley	33
4.1	The state of algorithms, now	33
4.1.1	Who reserves judgement?	33
4.1.2	What do we do with the replaced?	35
4.1.3	Privacy in the age of AI	37
4.2	How close are they, really?	37
4.2.1	What will the use us for?	39
4.2.2	Opportunity and risk are indeed twins: with new opportuni- ties come new problems.	41
4.3	Regulation	41
4.4	The control problem	42
4.5	Global reactions	43
4.6	What's next?	45
4.6.1	We've been at crossroads before	46
5	Conclusion	49
A	References	51
A.1	Papers	51
A.2	Books	56
A.3	Reports	56
B	Derivations	59
B.0.1	Derivations for the Basic Model	59

List of Figures

1	New Yorker Magazine, August 24th 2018. Artist: Farley Katz.	v
3.1	An example of a map of the current state of algorithms as they map to AI. Source: Oracle report on chatbots (2017)	16
3.2	A comment on Machine Learning...	18
3.3	An example of image classification by a deep convolutional neural network, from LeCun et al., 2015	19
3.4	An example of how Natural Language Processing analyzes text.	22
3.5	A Deep CNN layout for classifying skin lesions from Esteva et al., 2017	23
3.6	From the Australian megatrend report on "Tomorrow's digitally enabled workforce." Reeson et al., 2016	27
4.1	From Masahiro Mori's work in 1970 which coined the term the "Uncanny Valley" - bukimi no tani genshō	34
4.2	The tendency towards disruption by algorithms. From Wright & Schultz, 2018	36
4.3	Tricking a clever algorithm. From Nguyen et al., 2015	40
4.4	Global reponses to AI. Dutton, (2018)	43

Dedicated to my daughter Vera W. Mohn...

Chapter 1

Why tasks? Why algorithms? Why work at all?

If natural phenomena have an air of "necessity" about them in their subservience to natural law, artificial phenomena have an air of "contingency" in their malleability by environment.

Herbert Simon

1.1 A short introduction

The robots are not coming for us. Not yet. However, algorithms will be the first widespread cognitive augmentation of the human intellect¹. It will usher in a new era of automation.

Before the industrial revolution, there were artisans (Richardson, 2008). During the industrial revolution, there were Luddites; and the Luddites were right, the machines did take their jobs (O'Rourke et al., 2013)². In fact, it took around sixty years before the benefits of the industrial revolution finally reached the average worker in the United Kingdom (Feinstein, 1998).³ That is a long time. Korinek & Stiglitz

¹Although the internet gave us the information and communication superhighway, in how we harness algorithms they will analyze, sort, process, and form information for us.

²Hobsbawm (1952, p.59) terms this "collective bargaining by riot.", claiming it was not about hostility towards machines but the consequence of the machines to their livelihoods; a technique of trade unionism.

³When discussing the industrial revolution in the United Kingdom, it is important to place events in historical context. There were many simultaneous difficult circumstances: an unprecedented doubling of the population of England and Wales (Chambers & Mingay, 1966), a succession of harvest failures caused by abnormal weather conditions (Chambers & Mingay, 1966), not to mention wars with France which withdrew thousands of men from civilian labour (Mokyr & Savin, 1976), created inflationary financial conditions (Mokyr & Savin, 1976), and stimulated some sectors while disrupting others (Mokyr & Savin, 1976). Some could compare these to similarly difficult circumstances faced by the world today.

(2017) are concerned with saving humanity from Malthusian destiny. Frey & Osborne (2013) have dire warnings about 47 percent of current job types. So will it take sixty years for us to adjust to algorithms? Perhaps algorithms are merely the evolution of digitalization, if the first step was the widespread implementation of personal computers and the huge ramifications that combined with immediate access to communication and information through the world wide web, the internet. Although displacement is a serious concern for labour worldwide, it is not the only answer.

What the world is facing now is not a world where a robot surgeon performs heart surgery, replacing the human surgeon. It rather is the augmentation of the human surgeon through computer-enhanced instruments which provide "superhuman dexterity through tremor filtration and motion scaling that are capable of precise manipulation in confined body cavities" (Zenati, 2001, p.1). Both physical and cognitive augmentation is the true power of the algorithmic revolution. This augmentation is here presented as best understood through a task-based model.

An algorithm is an ordered and finite set of operations that must be followed in order to solve a problem (Gersting, & Schneider 1995). Its purpose is to solve a problem, which means that it has a defined objective. So, what are algorithms doing? For example, algorithms are completing tasks. These tasks are the smallest identifiable essential piece that serves as a unit of work. A job is then a series of tasks that a worker performs. These tasks may be the same or they may evolve over time. This perspective differs from the technology augmenting perspective of "skills", where a skill is "a worker's endowment of capabilities for performing various tasks" (Açemoğlu & Autor, 2011, p. 1045).

Workers apply their skills to tasks and produce outputs in exchange for wages. In the past, the endowment of skills was how workers were judged, with high and low skilled workers as different cases in task-based models (Açemoğlu & Autor, 2011). By moving this logic from skills to tasks, a more highly skilled worker becomes a worker that can perform a higher variety, or range, of tasks. These tasks can change in response to labour conditions and augmentation by technology. A worker that can adjust to working alongside algorithms is the particular case that I will propose in this investigation.

This work proceeds as follows: in chapter two I discuss at length a model of task-based automation by Açemoğlu & Restrepo, their interpretation of the problem, and my own interpretation as well. I also discuss in part the model of task-based automation as presented by Korinek & Stiglitz, and their conclusions. Naturally, I add my own.

In chapter three I discuss algorithms at length. I define important algorithms present in discourse today, how they are used, how they are combined, and the implications

to the labour force, to local, and to global economies of this shift. I also discuss empirics of what is happening now, and how this process is accelerating.

In chapter four I present policy observations and policy recommendations for different nations, both developed and developing, for handling a new, "algorithmic" future. I also discuss the shortcomings of algorithms and of narrow intelligences, and how human labour will remain relevant, if not crucial, for decades to come, especially in the realm of working with algorithms.

In the conclusion I pull all these parts together around my central argument: that we will work with algorithms, and that they will revolutionize the future of work.

1.2 Working papers

I have conducted an extensive literature review on the economics of algorithms, artificial intelligence (AI), robotics, the implications of a world economy, and the labour market. The generation affected by a potentially massive change is quickly growing and often exploratory. Although published works exist, and are referred to, they are dated. The subject moves quickly. Ideas move quickly. A canonical model has yet to be formed, although it looks to be the idea of the task-based automation model. I acknowledge that this could be considered a weakness in my work and ask for it to be considered a strength instead. Great minds are working on this. They are processing their ideas through models, philosophies, ideas, implications, consequences, and more, even if John Maynard Keynes or Herbert Simon or Alan Turing already thought a little about it first.

Chapter 2

A model of task based automation

...we may say that the artificial object imitates the real by turning the same face to the outer system, by adapting, relative to the same goals, to comparable ranges of external tasks.

Herbert Simon

In this chapter, I outline and discuss a task-based automation model proposed in Restrepo and Açemoglu's (2018) working paper. This model uses a Cobb-Douglas production function with machines and labour as the two inputs, where different tasks are conducted by different inputs. Included in my discussion and critique of the model are the changes I propose to better understand the world where humans work by harnessing algorithms. The core argument is that there will be multitudes of tasks that can only be completed with the combination of algorithms and labour. This production function would be an interesting measure of the degree of automation in an economy, and perhaps form the basis with which you could compare and contrast different nations' levels of automation and adoption of algorithms. After all, human beings and algorithms working together can create an entirely new future. If the human mind is the ultimate resource, combining the mind with ultimate narrow intelligence sounds like we are strapping ourselves into a mental accelerator of fantastic proportions.

2.0.1 On algorithms

An algorithm is defined as a separate input to labour, as an algorithm is not entirely labour replacing. As defined earlier, an algorithm is an ordered and finite set of operations that must be followed in order to solve a problem (Gersting, & Schneider 1995). Its purpose is to solve a problem, which means that it has a defined objective. Algorithms and labour are not interchangeable factors, as most algorithms need labour to work with or alongside it (Seamans & Raj, 2018). For example, a machine learning algorithm that uses natural language processing to read contracts

cannot also find clients, nor can it participate in meetings to convince people to hire the said algorithm (Alarie et al., 2017). For that, you would need a human being working with an algorithm, as it can complete certain aspects of a profession, that is to take on tasks. It could be that by taking on these tasks or opening the opportunity to creating brand-new tasks, this may redefine the profession. Much like dishwashers that do not wash dishes in the way that human beings wash them, algorithms perform tasks in an entirely different manner to human labour. Thus, a model of task-based automation a good starting point to how the widespread adoption of algorithms in advanced, and even less advanced, countries' economies will affect the world economy. This does raise the question of displacement of human labour.

2.0.2 Displacement?

In the past, a job was considered in its whole (even if there are heterogeneities in job titles, see Frey & Osborne, 2013). However, with the changes in the labour market, the rise of skill-sharing, of the gig economy, there has been in a shift in modelling from skills to tasks, where these parts described as tasks are best understood through a task-based automation model. A task is the smallest identifiable essential pieces that serve as a unit of work (Açemoğlu & Autor, 2011).

If we shift understanding from skills to tasks, then a person's displacement in the job market is not absolute. The person is merely nudged to one side, in some aspects. This notion creates what Açemoğlu and Restrepo (2018), p.1 term the displacement effect "... a decoupling of wages and output per worker, and a decline in the share of labor in national income".

If this displacement effect does not only unlink wages and output per worker but then also lead to a decline in the share that workers have in national income, this has important implications for inequality. The authors Korinek & Stiglitz (2017) address this in part, which will be covered later in the chapter.

So although automation, algorithms, and AI are productivity enhancing, they can also be in part labour-replacing, especially when viewed from the perspective of tasks rather than skills, as modelled here with task-based automation. To finish, these ideas, like many large notions, are not new to great minds. To quote Wassily Leontief: "Any worker who now performs his task by following specific instructions can, in principle, be replaced by a machine. This means that the role of humans as the most important factor of production is bound to diminish—in the same way that the role of horses in agricultural production was first diminished and then eliminated by the introduction of tractors." (Leontief, 1983, p.3)

2.1 Restrepo & Açemoğlu's model, with comments

Restrepo and Açemoğlu's model begins with an equation:

$$\ln Y = \int_{N-1}^N \ln y(x) dx \quad (2.1)$$

which details an economy where the output is an aggregate function Y of the tasks x . This type of framework appeared in an earlier model of Açemoğlu & Autor (2011) where they considered a static environment with a single final good, a closed economy, and no trade in tasks (itself an interesting concept). This final good is produced by combining a continuum of tasks as represented by a unit interval $\in [0, 1]$.

$$Y = \exp \left[\int_0^1 \ln y(i) di \right] \quad (2.2)$$

where Y denotes the output of a unique final good, with $y(i)$ the production level of task i . Constructing a model as such shows output, or final goods, as the product of tasks performed by workers of varying capacity. Açemoğlu & Autor (2011) model only have workers performing tasks, while Restrepo and Açemoğlu (2018) also have "machines" aka capital performing tasks. The next iteration I would recommend is to consider the combination of labour and capital, or how algorithm can indeed augment human labour, and to consider this an individual input. That production function in itself would be an interesting measure of how automated an economy is becoming, for example to compare what continuum or range of tasks is taken on by only machines, by only humans, and by the combination of algorithms and humans.

Looking deeper into the construction of the model, Restrepo and Açemoğlu (2018) also refer to Zeira (1998). In Zeira's paper from 20 years ago (!), he has a single good Y which is used for both consumption and investment. This good is produced by an intermediate good X , produced by labour and capital.

$$Y = aX \quad (2.3)$$

Zeira's 1998 paper considers productivity differences across different nations, and thus he introduces the augmenting parameter a which is the productivity parameter that may differ across various nations around the world. He considers two types of inputs, one low technology with low level labour, and one high technology with high level labour. Thus, the model of task-based automation as presented in the latest paper by Restrepo and Açemoğlu draws upon a model built to compare differences both within a nation, and between nations. That kind of basis then lends itself well to both discussion about augmentation of labour by algorithms (for example, the

productivity measure a on the intermediate good X could be modelled as a measure of the degree of adoption of algorithms in different nations).

Zeira's 1998 paper is also curiously relevant in that it details the productivity differences in countries with differing levels of capital. This framework of thinking is informative to the current paper's model in terms of determining who will own the algorithms, if algorithms are considered to be capital rather than technology or if algorithms are something that can be owned.¹

Going back to Restrepo and Açemoğlu (2018), they proceed to define a task x as within the range:

$$x \in [N - 1, N] \quad (2.4)$$

Defining it as such frees the authors from the real number range imposed in the original model by Açemoğlu and Autor (2011) whereby:

$$i \in [0, 1] \quad (2.5)$$

that is, that the continuum of tasks belonged to a limited range between 0 and 1. As Restrepo and Açemoğlu (2018) acknowledge in their paper, adjusting the range gives them the freedom to add or remove tasks to the continuum without needing to alter the range of tasks from which the production function draws.

But what is a task? If x denotes a task, this is considered the smallest identifiable essential piece that can be considered a unit of work. A task could be sending an email, finding statistics for a report, laying a tile, reading a contract, finding someone's medical records, and a profession or a job is made up of a series of tasks. Tasks are the smallest building blocks of an economy, perhaps these blocks that fit together to form different professions.

In Restrepo and Açemoğlu (2018) each task can be produced by human labor, $l(x)$, or by machines, $m(x)$. The productivity of the type of input is represented by $\gamma_L(x)$ for labour or $\gamma_M(x)$ for machines.

$$y(x) = \begin{cases} \gamma_L(x)l(x) + \gamma_M(x)m(x) & \text{if } x \in [0, I] \\ \gamma_L(x)l(x) & \text{if } x \in [I, N] \end{cases} \quad (2.6)$$

¹There is much to be said here about cryptocurrency. Previously, capital was "controlled" by the government, in the sense that the government ultimately decided how much money there was in an economy. With the rise of cryptocurrency, the 'amount' is not set in the same way. This is a similar concept to the ownership of algorithms. Will they be owned by governments? Corporations? Open-source? A combination? How will you buy and trade them? Do they depreciate like normal capital? Can developing countries obtain a unique foothold if the currency for creating algorithms is not initially money, but minds? Answering these questions could be an entire body of work on its own. However, this is beyond the purpose of the current body of work.

Productivity is output per hour of a task (Autor & Salomons, 2018). There are varied outputs per hour of the individual factors and of the combined factor of production. This is explained by labour performing tasks differently to algorithms, and again the combination of labour and algorithms performing tasks in another manner again than just labour would or just algorithms would. The tried and tested example of how a dishwasher washes dishes differently to a human being is an analogy to what informs this distinction.

Restrepo and Açemoğlu (2018) assume that productivity of labour, $\gamma L(x)$ is increasing and that the ratio between the productivity of labour and machines is increasing in (x) , thus giving labour a comparative advantage in higher-indexed tasks. The authors note that their framework builds on Zeira(1998) where he models how the difference of the rate of adoption of capital-intensive production methods amplify productivity differences across different nations and across time. Restrepo and Açemoğlu (2018) build on this logic to show that the rate of adoption of what they term machines, which is what I term algorithms, could result in four different types of technological advances, and how these advances impact the demand for labour and wages.

Although this framework is valuable in the discussion of the either/or situation, it misses out on discussing what I believe is the most valuable and informative case: the case of algorithms necessarily working with humans. I believe there are tasks, and perhaps in the future entire professions, that depend on the combination of labour and algorithms. This can be observed now with the necessary skill of using a personal computer for working in most middle-class, white-collar roles. I propose that algorithms will play this role in the future, that to work in a world with algorithmic labour, human labour will work alongside and necessarily with algorithms. I also propose that this case, for example it could be termed:

$$\gamma z(x)z(x) \quad \text{if } x \in (I, F) \quad (2.7)$$

The range of tasks $x \in (I, F)$ would lie between the previous two cases. Thus, on the lower end of the continuum would lie tasks best performed or best suitable to algorithms, in the middle is the combination, and on the end are tasks that can only be performed by humans. The size of these ranges as relative to one another would give insight then into the level of automation of an economy, and perhaps those tasks that are performed by both humans and algorithms could show the healthy side of automation, where humans are not only being replaced. The combination of human and algorithm tasks are placed in the middle of the task continuum to reflect an idea that tasks that can be exclusively performed by labour without algorithmic presence will be highly specialized, equivalent to higher productivity or higher wage tasks, and that the wage paid to labour in labour only tasks $W_H > W_Z$. Labour in this

model is supplied in the case of with and without algorithms, and thus there are two expressions for labour.

In Restrepo and Aćemođlu (2018) the assumption is that tasks now performed by algorithms (machines) could be performed by labour as well, in the sense that complex calculations were possible before Matlab, just much more tedious. Certainly, by freeing economists from tedious statistical calculations using advanced econometric methods that can harness machine learning, for example, more or different data can be analyzed. Perhaps this could be an argument for why the continuum perhaps should be sorted with algorithms and labour combined being on the higher side.

In their model, Restrepo and Aćemođlu (2018) denote the equilibrium wage rate by W and the equilibrium cost of machines by R , familiar terms for the price of labour and capital. Their equilibrium prescribes that firms choose the cost-minimization method of producing each task, and that labour and machine markets clear. Thus:

$$\frac{\gamma_L(N)}{\gamma_M(N-1)} > \frac{W}{R} > \frac{\gamma_L(I)}{\gamma_M(I)} \quad (2.8)$$

The second inequality implies that all tasks in $[N, I]$ will be produced by machines. The first inequality implies that the introduction of new tasks—an increase in N —will increase aggregate output. This assumption is imposed on the wage to rental rate ratio, which is an endogenous object. According to the empirical work performed by Autor & Salomons (2018) following this model, the assumption of wage ratios has no bearing on reality as their analysis is silent on new task creation. However, this assumption also implies that capital is more cost-effective than labor in newly automated tasks, which would mean tasks could tend more to the automated.

$$Y = B \left(\frac{K}{I - N + 1} \right)^{I - N + 1} \left(\frac{L}{N - F} \right)^{N - I} \quad (2.9)$$

Then output Y in equilibrium is a product of the different relative productivities and the range of tasks that the different inputs of production cover.

$$B = \exp \left(\int_{N-1}^I \ln \gamma_M d(x) + \int_I^N \ln \gamma_L d(x) \right) \quad (2.10)$$

B is a collection of terms that encompass productivities of the different factor inputs. As commented by Autor & Salomons, conventionally this result corresponds to Total Factor Productivity (TFP), which can then shift depending on if one or both of the efficiency terms $\gamma_M(x)$ or $\gamma_L(x)$ change, or because tasks are reallocated from one factor input to another. As the authors note, as distinct from the canonical Solow model, TFP growth in this setting is not Hicks-neutral if it stems from movements in either I or N .

Through this equilibrium result, by taking derivatives at the margins, Restrepo and Açemoğlu (2018) map out the displacement effect but argue that this is counteracted by the productivity effect which would push towards increased labour demand.

$$d\ln W = \underbrace{\frac{d\ln(N-I)}{dI}}_{\text{Displacement effect } < 0} + \underbrace{\frac{d\ln(Y/L)}{dI}}_{\text{Productivity effect } > 0} \quad (2.11)$$

Restrepo and Açemoğlu (2018) claim that without the productivity effect, automation would always reduce labor demand, because it is directly replacing labor in tasks that were previously performed by workers. They argue that if the productivity effect is limited, automation will reduce labor demand and wages. Autor & Salomons (2018) take this analysis further to suggest two channels through which this productivity effect may operate. The first, more direct effect is that automation may increase labour demand in non-automated tasks in an industry where automation is taking place. They call the 'Uber' effect, i.e., a technological improvement that both raises labor productivity and employment in the affected sector. Another productivity effect, this one considered indirect, would be that productivity growth in a technologically advancing industry may raise labour demand in other industries through rising consumer incomes and boosting final demand. They term this the 'Walmart' effect as automation lowers input costs to downstream customer industries, leading to output and employment growth.

Finally, Autor & Salomons (2018) discuss an avenue by which technological change may affect output and wages: through the creation of new tasks in which labor has comparative advantage. that is, a rise in N . An important distinction, and certainly a method for which labour can argue its relevance.

However, as noted earlier, alongside some equations in Appendix B, what this model and other interpretations or empirical work performed using this framework lack is an understanding of the necessary combination and possibilities of humans working with algorithms. Much like this master's thesis could not be completed without access to the internet or access to a computer, algorithms will become a necessary part of the future workforce. The question is how to harness them, and how to model this appropriately, and not just through the lens of either or. Displacement is not the only answer.

2.1.1 Korinek & Stiglitz

Korinek & Stiglitz, 2017, model task-based automation in a different manner. They are especially concerned with whether or not innovation in the algorithm space, or as they call it, AI, will be labour-augmenting or labour-saving. They link this to a given wage and flesh out a model around that.

The authors discuss that advances in AI could be considered a continuing evolution of a long process of automation that started with the industrial revolution. They also note others, myself included, who emphasize that AI critically differs from past inventions. We have yet to invent something that could make us completely obsolete. James Barrat (2013) calls this our final invention.

They also address an important idea: that productivity-enhancing technologies have previously always increased overall labour demand, not to mention how labour's share of national income has remained roughly constant (Jones and Romer, 2010), a notion even Keynes himself (1939) deemed a bit of a miracle. Algorithms, robotics, and AI could be the first case to overwhelmingly prove that wrong (Autor & Salomons, 2018). However, they don't consider how human beings will work with algorithms, only the "instead of" case. I believe this is a misnomer and a misunderstanding of how freeing human labour and combining it with the processing, calculating, predictive, and analytical power of different algorithms could be the ultimate productivity enhancer.

However, this is not reflected in the empirics (Autor & Salomons, 2018). Measured productivity seems to have increased rather slowly in recent years. If innovations related to algorithms and artificial intelligence enter the economy at the same slow pace as suggested by recent productivity statistics, then the transition will potentially be even slower than the wave of mechanization in the 1950 – 1970s (Korinek Stiglitz, 2017). I disagree with the authors that then this means the disruptions may not be that significant. Indeed, empirics as discussed later indicate the opposite.

The main argument in Korinek Stiglitz, (2017) revolves around Pareto-improving transfers. They suggest, for example, changes in patent length and capital taxation which could act as a device to redistribute to displaced workers. Their first best alternative is lump sum transfers between displaced workers and the innovators behind the algorithms. Certainly, there is the question of who could afford the goods and services created by algorithms and robots in a future where many workers are displaced.

For Stiglitz & Korinek (2017), for example, they outline how in the short run, the additional unit of machine labour added to an economy would earn its marginal product while simultaneously redistributing income from labour to traditional capital in that it alters the ratio between the two. This of course is not taking into consideration my main point throughout this thesis that algorithms and human labour will necessarily work better together, and to a degree need each other.

What Korinek Stiglitz, (2017) offer is a different perspective on task-based automation and its implications on the future labour market, which combined with empirics from Autor & Salomons, (2018) suggest that the potential productivity increase from the increased adoption of automation and algorithms is yet to appear.

2.1.2 A note on capital: will algorithms be considered capital, labour, or technology?

Often in these models, algorithms are defined as capital in the sense that they could be considered machines, or could be bought and sold, or are owned by companies, which to a degree makes sense when trying to model a fairly new idea into strict concepts of the factors of production. However, I would argue that algorithms are unlike previous factors of production. Although companies can own algorithms, algorithms can also be substituted for human labour at least in the sense of certain tasks, a range which is only growing. Unlike a machine that perhaps replaced a human being on a car production line, an algorithm could learn and grow, and create entirely new possibilities. It is almost like a capital that could invest in itself. Perhaps human capital is another lens through which to perceive algorithms. In any case, algorithms are certainly unusual. As is the definition of capital.

Capital has been a tenuous question in economics before. The great Joan Robinson in her 1953 treatise on *The Production Function and the Theory of Capital* raises many important questions about the notion of capital, about the particular list of objects it can entail, and about how different individuals may perceive those objects as different even if the good has the same type of description. But her question of capital in terms of cost or value is what I feel is relevant here. How will we value algorithms? What do they cost? What will they cost us?

Chapter 3

The What and the How:

If they find a parrot who could answer to everything, I would claim it to be an intelligent being without hesitation.

Denis Diderot

3.1 What we talk about, when we talk about algorithms

Terms like AI, algorithms, machine learning, and natural language processing are littered throughout this work. Consequently, some clear definitions of what these terms are, and how they are used, will now be presented.

Many narrow intelligences, or algorithms, are inspired by the human nervous system and the human mind. Despite this, they operate quite differently than humans. It may become a tired example, but the logic is solid: much like how a dishwasher does not wash dishes like a human being, a deep learning algorithm cannot tell the difference between a Samoyed and a Pomeranian in an image like a person would (Le Cun et al., 2015).

Much like the dishwasher and other modern appliances revolutionized everyday life at home, "...twenty-first century AI enables a constellation of mainstream technologies that are having a substantial impact on everyday lives" (LeCun et al., 2015, p. 436). From asking a speaker such as Amazon Alexa or Google Home or on an iPhone, Siri, what the weather is like today (and getting a correct answer), to self-driving automobiles like the Tesla model X, to robots that perform backflips (Boston Dynamics), they are becoming a part of our mundane reality.

3.1.1 Artificial Intelligence

Throughout this work, I talk about AI through the lens of many narrow, specialized intelligences (e.g., machine learning, neural language processing, deep learning, image recognition). This is because AI is not a generalized type of intelligence. There

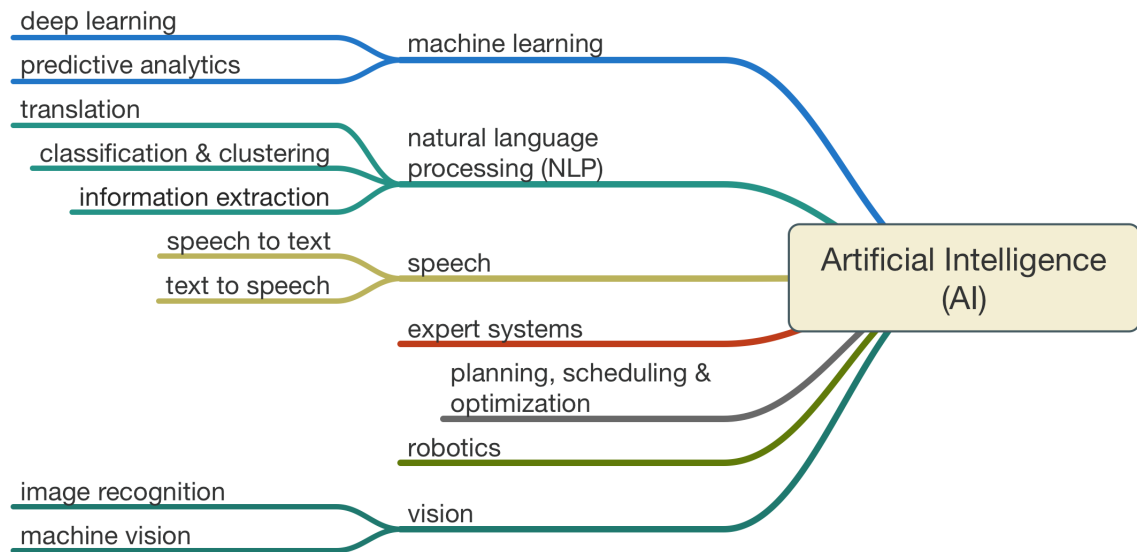


FIGURE 3.1: An example of a map of the current state of algorithms as they map to AI. Source: Oracle report on chatbots (2017)

is no AI that compares to the human mind; in AI terms, this is called Artificial General Intelligence (AGI) (Goertzel, 2007).

For AGI to exist, I agree with Monica Anderson, who believes that AGI cannot be created through models, or through reductions. She argues that current, narrow AI is much like Kahneman's System Two of thinking. She proposes that true AGI lies in creating something that is more alike Kahneman's System One: Intuitive AI ¹

What is the power of AI? Agrawal et al. (2017) argue that it is a predictive machine like no other. However, this seems to be a limiting concepts of the true capabilities of artificial intelligence, or algorithms. After all, human intelligence is not merely predicated on the ability to predict. It also includes the ability to create, to synthesize, to empathize, to understand, to learn, to grow, to process, to move, perhaps even to love and to forgive.

3.1.2 Machine Learning

Machine learning is a field that develops algorithms for predictions, classifications, and the clustering or grouping of tasks. Machine Learning involves training algorithms on large data sets to find patterns, programming computers to optimise a performance criterion using example data or past experience. (Alpaydin, 2010). There are two main branches of machine learning: supervised and unsupervised.

A supervised machine learning algorithm uses set inputs, for example features or covariates to predict an outcome (Athey, 2018). As noted in a somewhat snarky image later, when it comes to the term prediction for machine learning, it is not

¹See Monica Anderson's rather amazing array of work. This topic is sadly much too broad to fit into this body of work.

about forecasting, but rather about using training data to "teach" the algorithm to find the predicted outcome.

There are many applications of machine learning taking on tasks (e.g., reviewing legal documents). The Heretik website (i.e., <https://www.heretik.com>) is an example of just this. More specifically, "Heretik is a relativity application that marries effective text analysis machine learning models and flexible document review capabilities."

Machine learning is not a method that is necessarily new to econometrics. Varian (2014) discusses how econometric methods have much to gain from machine learning, for example with prediction. As he notes on page five, "Econometricians, statisticians, and data mining specialists are generally looking for insights that can be extracted from the data. Machine learning specialists are often primarily concerned with developing high-performance computer systems that can provide useful predictions in the presence of challenging computational constraints."

Machine learning is closely related to data mining, a method employed by the Australian Taxation Office (ATO) to detect fraudulent tax statements (ATO, 2015).

The other field of machine learning is called unsupervised ML, which involves finding clusters of observations that are similar in terms of their covariates. For example, if a computer scientist used an unsupervised ML algorithm to categorize what content appeared in videos, and when a human being judged the largest group of videos, they determined that most of the videos contained a specific type of content, for example cat videos. This is referred to as "unsupervised" because there were no "labels" on any of the images in the input data; only after examining the items in each group does an observer determine that the algorithm found cats or dogs (Athey, 2018).

In my view, these tools are very useful as an intermediate step in empirical work in economics. They provide a data-driven way to find similar newspaper articles, restaurant reviews, etc., and thus create variables that can be used in economic analyses."

An example of ML in use is by the fantastic paper by Mann & Puttman on the benign effects of automation. Puttman wrote an algorithm² to classify all five million U.S. patents granted between 1976 and 2014 as automation or non-automation patents. The authors documented a rise in the share of automation patents from 25 percent to 67 percent. They find a positive effect of automation on total employment, driven by job growth in the service sector, which compensates for a fall in manufacturing employment (Mann & Puttman, 2017). What a fantastic example of human beings working with algorithms!

²I had ambitions of requesting an algorithm to read papers for this investigation, thus proving the use of ML in economic research. However, I ran out of time. I look forward to using those techniques in the future.

Interviewer: What's your biggest strength?

Me: I'm an expert in machine learning.

Interviewer: What's $9 + 10$?

Me: Its 3.

Interviewer: Not even close. It's 19.

Me: It's 16.

Interviewer: Wrong. Its still 19.

Me: It's 18.

Interviewer: No, it's 19.

Me: it's 19.

Interviewer: You're hired

FIGURE 3.2: A comment on Machine Learning...

3.1.3 Deep Learning

Deep learning is a field of machine learning which builds upon the techniques described above to a greater layer of complexity through "nets". It is often termed "convolutional neural networks" To quote one of the giants of the Artificial Intelligence field, Geoffrey Hinton:

"Deep learning allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction...Deep learning discovers intricate structure in large data sets by using the backpropagation algorithm to indicate how a machine should change its internal parameters that are used to compute the representation in each layer from the representation in the previous layer." (LeCun et al., 2015, p4)

The deep in deep learning refers to the hidden layers present in the neural network. "Deep learning refers to the strategy of using architectures with many hidden layers to tackle difficult problems, including vision."(Kriegeskorte, 2015, pp.4).

DL methods have dramatically improved a number of fields previously considered very challenging to artificial intelligence: speech recognition, visual object recognition, object detection (Esteva et al., 2017). Deep convolutional nets have brought about breakthroughs in processing images, video, speech and audio.

Deep Learning is often applied in conjunction with Machine Learning or Natural Language Processing in fields by various startups in the legal tech space, that is where entrepreneurs create algorithms to read contracts, write contracts, read law, analyze cases, and other tasks previously considered grunt work or taken on by law interns or paralegals (Alarie et al., 2015).

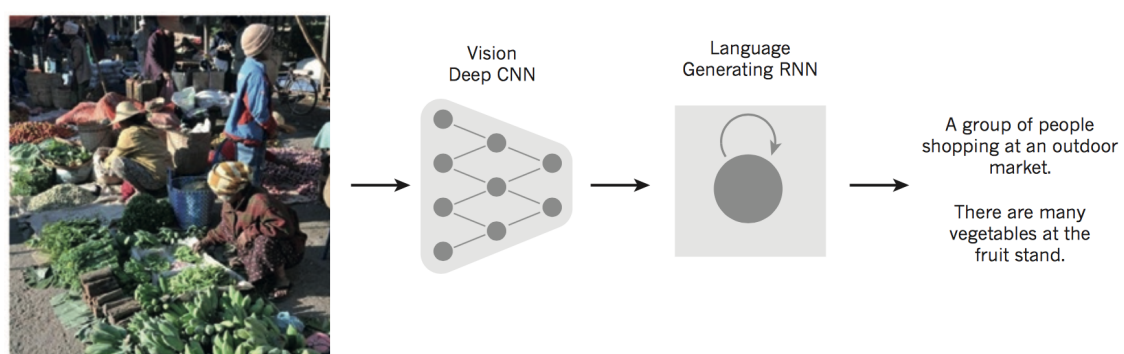


FIGURE 3.3: An example of image classification by a deep convolutional neural network, from LeCun et al., 2015

3.1.4 Reinforcement Learning

Reinforcement learning (RL) is a framework that shifts the focus of machine learning from pattern recognition to experience-driven sequential decision-making. It is a set

of methods that allows a machine to be fed with raw data and to automatically discover the representations needed for detection or classification. RL is in fact familiar to economists, for it roughly lines up with dynamic programming (Athey, 2018). It promises to carry AI applications forward toward taking actions in the real world (Stone et al., 2016).

RL could assist in providing a normative account of how agents may optimize their control of an environment. An early version RL has been used by Game Theorists within economics to map strategic decisions (Erev & Roth, 1998). To use RL in situations approaching a semblance of real-world complexity, this presents a difficult task as agents must derive representations of the environment from their high-dimensional sensory inputs, and then use these to generalize past experience to new situations.(Mnih et al., 2015). A great weakness of algorithms is generalizing any kind of information. Algorithms are narrow intelligences. Human beings and the human mind provides the generalization.

3.1.5 A comment on the term "learning"

Many algorithms call themselves "learning": Machine Learning, Deep Learning, Reinforcement Learning, however these are not "learning" in the true, human sense of the word. With great respect paid to Alan Turing, perhaps the originator of ideas of Artificial Intelligence. His comment in his paper from 1953 removes the philosophy and the ambiguity of the concept thinking. When he proposes that we consider the question "Can machines think?", he rightfully so adjusted in the imitation game to what a machine may do given a specific situation and specific parameters. Algorithms do not generalize. Machines do not generalize. Algorithms, can however, predict, and this can be mistaken for generalization.

3.1.6 Natural Language Processing

Natural Language Processing (NLP) is an example of an algorithm that could also be called text analytics. NLP helps computers interpret, understand, and then use everyday human language and language patterns, by extracting high quality of information from text (Feldman, & Dagan, 1995) It breaks both speech and text down into shorter components and interprets these more manageable blocks to understand what each individual component means and how it contributes to the overall meaning (Liddy, 2001). Through NLP our AI can transcribe consultations, summarise clinical records and chat with users in a more natural, human way (Babylon, 2018). NLP has wide applications, from medicine to law to economics. Plenty of different law-tech startups claim that their NLP algorithms can summarize thousands of pages of legal documents in minutes. Although there are claims that it reduces the probability of error, of course this assumes that the original way that the application

was coded was correct and that the algorithm was trained with sufficient data and a sufficient variety of data (McGinnis et al., 2014).

NLP techniques are also called text mining. An example of text mining used in economics is by Siegel, 2018, who uses NLP to analyze financial reports of different medium and large corporations in Germany. The algorithm first prepares the text by tokenization, that is to divide the text into sentences and tokens (words). Then she codes what the algorithm is searching for: terms like development, or sections of text that describe observations or predictions of the future. Finally, the algorithm is trained to structure this text in a way that an economist would find the summaries valuable.

3.1.7 Computer Vision

Computer vision is an example of machine perception. With the advent of affordable, large-scale computing, the availability of large datasets, and refinements of neural network algorithms has led to dramatic improvements in performance on benchmark tasks. In fact, there are now some narrowly defined visual classification tasks that can be performed more accurately by computer vision than by human beings (Sankaranarayanan et al., 2018). Much current research is focused on automatic image and video captioning (Stone et al., 2015).

Computer vision is the type of algorithm that is looking to replace professions such as dermatologists or radiologists. In a paper by Esteva et al., 2017, they present a CNN which is trained using a dataset of over 129 000 images that document over 2000 different diseases. They test this algorithm's performance against 21 board-certified dermatologists using images that were confirmed later by biopsy. The CNN achieved performance on par with all tested experts, demonstrating that an algorithm in a specific context trained on specific data could display competence on par with dermatologists. The authors suggest that mobile devices can potentially extend the reach of dermatologists outside of the clinic, as most individuals in developed countries now have access to a smartphone, and thus this could democratize vital diagnostic care by making it accessible to more individuals³.

3.2 What drives these algorithms?

What drives the algorithms?

- data
- processing power

³It would be interesting to document how different highly skilled professionals experience algorithms encroaching upon their expertise and their vision for the future of their field

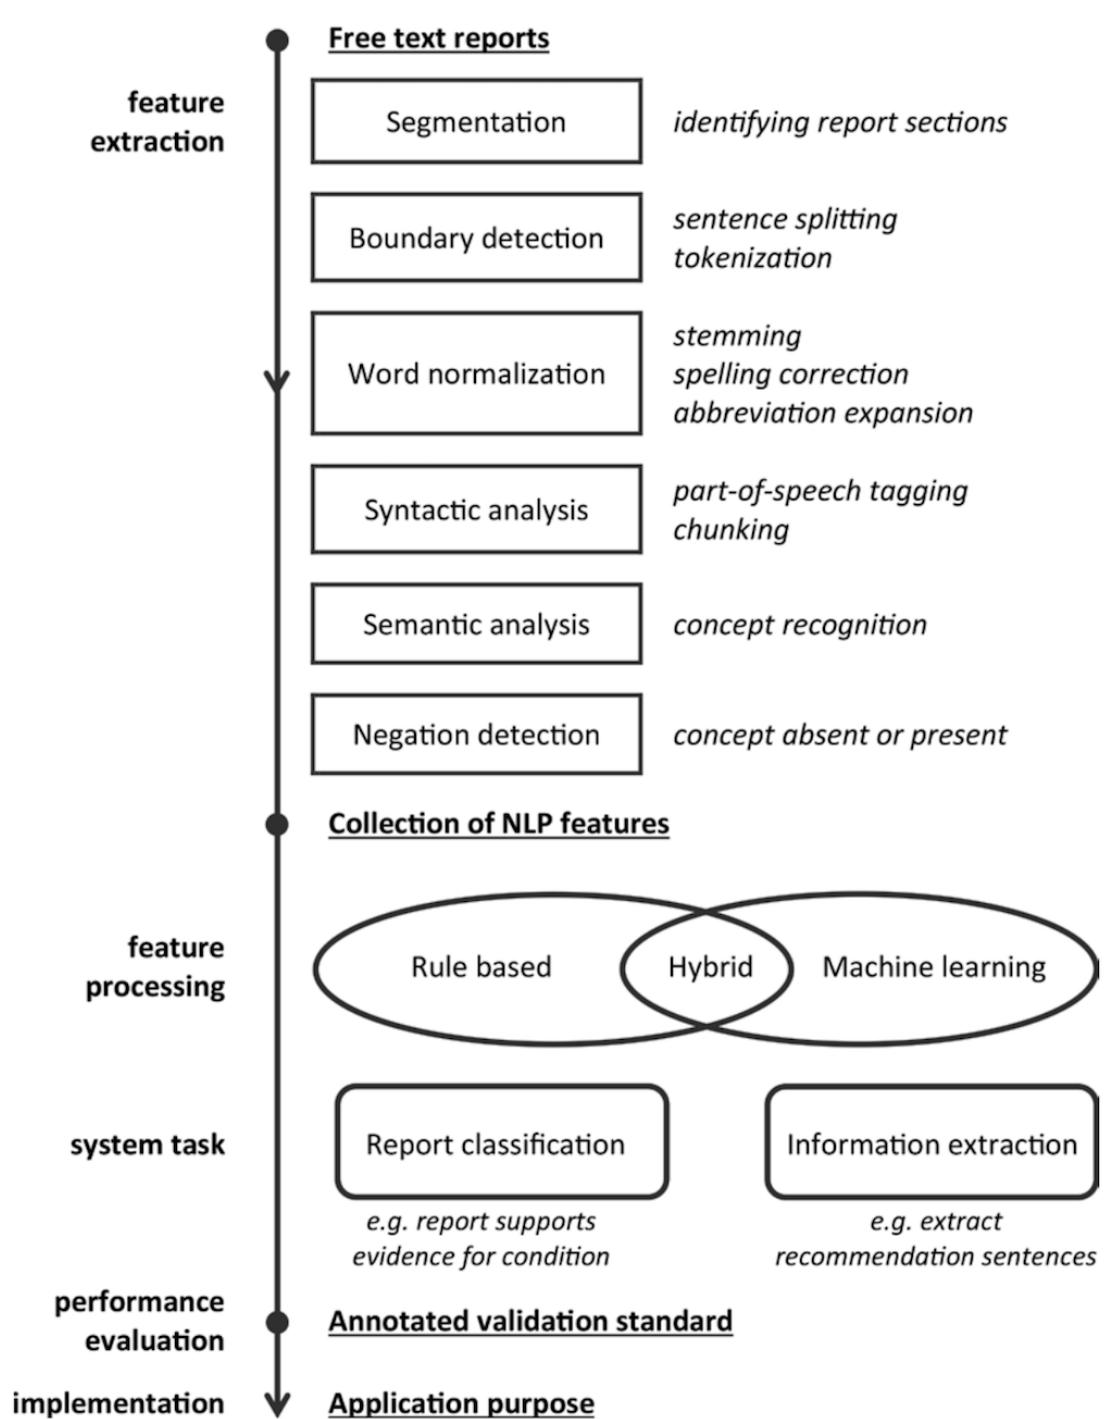


FIGURE 3.4: An example of how Natural Language Processing analyzes text.




Figure 1 | Deep CNN layout. Our classification technique is a deep CNN. Data flow is from left to right: an image of a skin lesion (for example, melanoma) is sequentially warped into a probability distribution over clinical classes of skin disease using Google Inception v3 CNN architecture pretrained on the ImageNet dataset (1.28 million images over 1,000 generic object classes) and fine-tuned on our own dataset of 129,450 skin lesions comprising 2,032 different diseases. The 757 training classes are defined using a novel taxonomy of skin disease and a partitioning algorithm that maps diseases into training classes

(for example, acrolentiginous melanoma, amelanotic melanoma, lentigo melanoma). Inference classes are more general and are composed of one or more training classes (for example, malignant melanocytic lesions—the class of melanomas). The probability of an inference class is calculated by summing the probabilities of the training classes according to taxonomy structure (see Methods). Inception v3 CNN architecture reprinted from <https://research.googleblog.com/2016/03/train-your-own-image-classifier-with.html>

FIGURE 3.5: A Deep CNN layout for classifying skin lesions from Esteva et al., 2017

- high-speed internet

Concepts behind the algorithms I have detailed above are not necessarily new, but what has triggered the explosion of progress to the degree that we are discussing a potential for replacement of human labour are the three factors outlined above: data, processing power, and high-speed internet. Access to these meant a wide variety of long-standing problems in ML, AI, and computer vision saw improvements to the extent that they broke through long-standing performance barriers (Lemley co., 2017).

Humankind generates and collects more data today than ever before (Domo, 2018). Data is generated by social media use, by appliances connected to the internet, by loyalty programs at stores, by exercise tracking applications on smartphones, by applications that track female fertility or menstrual cycles, by genealogy websites, and by internet history (Domo, 2018). There is a never-ending array of data points that are generated, collected, and potentially capitalized upon. The old adage of ‘if the product is free, what they are selling is you,’ only gets truer with time. Some researchers predict that the total amount of data generated annually by any device could reach 163 zettabytes (one zettabyte is equal to 270 bytes) by 2025 (Reinsel, Gantz, Rydning, 2017).

3.2.1 Disruption

Algorithms are disruptive. They are disruptive to how we think about capital, they are disruptive to how we think about labour, they are disrupting the future of work. The disruptive nature of algorithms and their combination of narrow intelligences into AI is difficult to structure neatly into a mathematical model. The movement feels at times equivalent to when the fish crawled out of the sea.

3.2.2 Demand for labour in the world of algorithms

There is a considerable fear around algorithms replacing labour. This is much of the focus of the work by many authors when discussing automation, algorithms, and artificial intelligence. The disruptions generated by depend on whether they are labor-augmenting or labor-saving, whether at a given wage, the innovations lead to more or less demand for labor (Korinek Stiglitz, 2017).

In modern, advanced economies (and to a lesser extent, but still relevant to less advanced economies) computers changed the workforce completely. Internet access has changed it again. Now, the speed and breadth of internet availability, the expected connectedness of co-workers, of managers, and the level of employment that is service based, and especially services which can be delivered remotely, implies that change can go in an entirely unimagined direction. It is difficult to see the possibilities before they exist, and difficult to predict the complexities of the upcoming changes.

3.2.3 Let's get hyper-specialized

Sometimes I feel that the age of the generalist is dead. A PhD, in many degrees, is a tiny, tiny piece of knowledge carved from a hill, placed on a pile, that is often never read again. PhD graduates become the queens and kings of the tiniest domains. The new professions created for, with, and alongside algorithms could reinvigorate the generalist. Instead of being valued for knowing specific knowledge (or knowing how to find it), if an algorithm reads papers and spits out information, what does the human being working with it do? They create! If the essence of the intelligence of the human mind, what is still difficult to mimic by an algorithm, is creativity, what better way to hypercreate than to have all the inputs prepared for you.

If some people code the algorithms, others test them, and algorithms test them some more, then there exists many new professions in this explosion of intelligence and creation that may well await us. That is, if people have access to the algorithms. I often wonder if corporations will replace governments. The question of the ownership of algorithms is a crucial one.

3.2.4 What about wages?

Wages are driven by supply and demand of the relevant labour force, whether skill based or task based, but also driven by demand. As noted by Brynjolfsson & Mitchell (2017) the elasticity of labor supply responds to wages. If there are many people who already have the requisite skills, then supply can be fairly elastic and wages may not fluctuate much, if at all, even in response to fluctuations in demand. In contrast, if

skills are more difficult to acquire, then changes in demand may be reflected more in wages rather than in employment (Brynjolfsson & Mitchell, 2017).

3.2.5 What about productivity?

Currently, the world is in an era of low productivity growth (Autor & Salomons, 2018). This is perhaps in spite of the only increasing adoption of algorithms. It certainly seems that we are on the cusp of widespread adoption of algorithms across many different fields, different professions, and different levels of cognitive ability. "We thus appear to be facing a redux of the Solow (1987) Paradox: we see transformative technologies everywhere but in the productivity statistics." (Brynjolfsson & Mitchell, 2018, pp.2)

3.2.6 What about job polarization?

Job polarization is a continuing trend whereby traditional middle-class jobs that require a middle level of skills are disappearing across developed economies (Goos and Manning, 2007). With the advent of algorithms, it is possible that this will be further exacerbated. (Açemoğlu & Autor, 2011). Job polarization contributes to inequality by broadening the wage gap between high and low skilled workers (Goos and Manning, 2007). It is possible that humans augmented with algorithms could present a new middle-class, middle-ground of the labour force. "As routine-based positions decline due to automation, middle-skilled workers will likely find themselves moving toward less skilled, service-oriented positions, leading to further wage erosion." (Wright & Schultz, 2018)

While these technologies may make society richer overall, unless the right governance structures are put in place, certain segments of society are likely to receive most of the benefits while others suffer most of the costs (Frey and Osborne, 2017). This is not the tide that lifts all boats, if it lowers the real wage and creates large swathes of unemployment.

3.2.7 Will we tax robots?

When considering job polarization, or mass unemployment, or the universal basic income, one answer people postulate is that there could be a redistribution of income from algorithms, or robots, to labour through taxation. Taxation of robots could discourage a full implementation of automation, and thus encourage companies to create work that human beings could do alongside, or with algorithms (Guerreiro et al., 2017). Korinek & Stiglitz, 2017, argue that since robots would be a fairly fixed

supply, that they could be taxed in a way that tended towards a Pareto improvement without creating distortions. This may also serve to fend off the Singularity, as mentioned later.

There is an estimated global shortfall of approximately 400,000 positions and this is forecast to almost quadruple by 2019. Retraining and education will be the key factors in filling the gaps created by different demands and new markets (Evans, 2017).

3.3 Are we headed towards Singularity ?

Singularity. The idea here is that rapid growth in computation and artificial intelligence will cross the boundar to dominate human labor (Korinek Stiglitz, 2017). Potentially economic growth could accelerate sharply as improvements cascade throughout the global economy (Brynjolfsson and McAfee, 2014). Certainly, looking at the control problem, we have created an invention that may continue to improve itself into perpetuity. This is an open question that is acknowledged here, but beyond the scope of this work.

3.4 Let's talk general Empirics

Automation is considered a considerable threat to the future of labour and the workforce. Frey and Osborne (2013) provides an idea of how large an impact automation could have on the U.S. labor force, and by generalization, to other nations as well. The authors focus particularly on ML as well as automation and propose a model to predict the extent of computerization's impact on non-routine tasks. After categorizing tasks by their susceptibility to automation, Frey and Osborne map these tasks to the O*NET job survey, which provides open-ended descriptions of skills and responsibilities involved in an occupation over time. They then integrate this dataset with employment and wage data from the Bureau of Labor Statistics (BLS) to determine what subsets of the labour market are at high, medium, or low risk of automation. The authors finds that 47 percent of U.S. employment is at high risk of automation.

Other authors have taken this methodology and applied it to individual nations, finding different ranges of the potential of automation. Arntz et. al, 2016, take into account the heterogeneity of workers' tasks within occupations, and find a much lower rate of risk of automation: only 9 percent of jobs in their study are automatable in the 21 OECD countries they chose, with ranges from six percent in Korea to twelve percent in Austria. Pajarinen and Rouvinen (2014) find that over a third of of Finnish jobs are at high risk of automation.

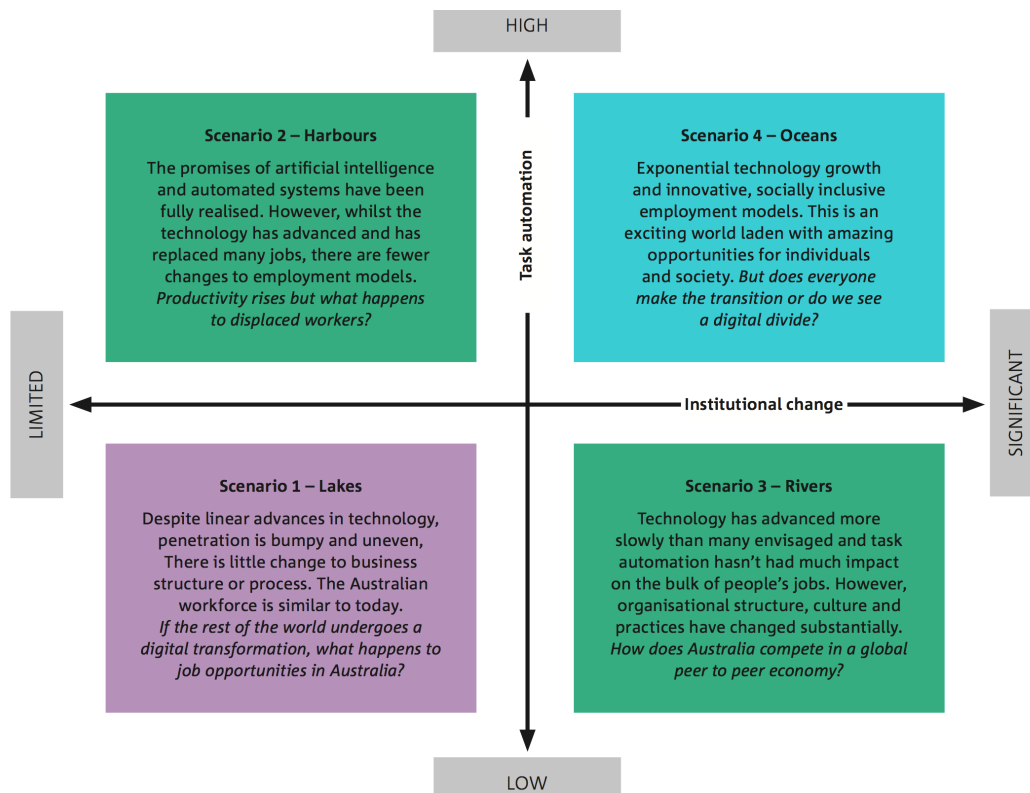


FIGURE 3.6: From the Australian megatrend report on "Tomorrow's digitally enabled workforce." Reeson et al., 2016

Taking these ideas of the risk of automation, I decided to examine three professions that are impactful when it comes to job creation and social mobility.

3.5 The effects of algorithms on three professions

In the following three sections I highlight the effect that algorithms have on three important fields: the medical profession, the financial profession, and the legal profession. The reason I chose these three is that they are often considered fields in which "anyone who works hard enough" can succeed, you merely need to study hard and work hard once you graduate. The issue that arises with the introduction of algorithms in these fields is that with firms taking on less graduates, chances are the graduates that are taken on are those that come from privileged families. I would consider these three fields as equalizers of a sort, for people who come from more 'salt of the earth' backgrounds that are looking to rise above their heritage. If these three fields no longer serve that function then this is an extra burden of inequality placed on society, not only the hollowing out of middle-class jobs as discussed by Osborne & Frey (2013), but the hollowing out of upper-class opportunities for those who are not born into that sphere.

3.5.1 The effect on the medical profession

It is unlikely that someone will walk into a doctor's office and be assisted by a robot. However, video-calling a doctor is an option available in many nations, including Norway. Babylon is a chatbot created to help diagnose patients before they see a GP in England, and is covered by the NHS (Babylon, 2018). These are tasks that are eroding from the purview of doctors, but perhaps in good timing, as doctors often seem very stretched for time.

The medical profession is ripe for augmentation. The earlier example of the "superhuman dexterity through tremor filtration and motion scaling that are capable of precise manipulation in confined body cavities" (Zenati, 2001, p.1) is important to note that with robotic assistance, surgeons can now accomplish what was previously impossible. Dermatologists could work with a deep CNN to confirm their diagnosis of skin cancer (Esteva et. al, 2017).

If algorithms can use NLP to read patients text or understand a recording of their voice, couple this with a ML algorithm trained on health data collected by that individual's smartphone (and many others), they could perform basic diagnostics, suggesting to the doctor who does see the patient a list of potential issues or symptoms and their potential diseases ranked by likelihood or possibility. This would be algorithms working with doctors. However, if algorithms end up replacing doctors, and doctors become more like customer service representatives (with wages adjusting accordingly), this amounts to displacement and leads to job polarization. The potential is there, but the direction is unknown.

3.5.2 The effect on the financial profession

Another profession famed for rags to riches (and perhaps rags again) is finance. With the advent of high frequency trading and algorithmic trading, stockbrokers and professionals in finance have changed. There is an increased demand for individuals who can write algorithms, however (Chaboud et al., 2014).

Although there are many who have positive views algorithmic trading, there are many different studies that imply that as with all things, the truth is multi-faceted and complicated. For example, Foucault et al., (2016) show that in a world with no asymmetric information, the speed advantage of algorithmic traders does not increase the basic purpose of prices which is to convey information, and rather it increases adverse selection costs. Another effect as shown by Jarrow and Protter (2011) is that algorithmic traders can end up collectively acting as one big trader, creating price momentum and thus causing prices to be less informationally efficient, another negative effect. Along these lines of acting at the same time, Kozhan and Wah Tham (2012) show that algorithmic traders entering the same trade at the same time causes a crowding effect, which in turn pushes prices further away from fundamentals.

The Securities and Exchange Commission (SEC) refers to HFT as a significant development of market structure in recent years that can be difficult to precisely define, that represent over half of all trades daily. It generally refers to professional traders engaging in strategies that generate a large number of trades on a daily basis (Menkveld, 2013). As outlined by a report issued by the SEC, the following are characteristics attributed to proprietary firms engaged in HFT:

- the use of extraordinarily high-speed and sophisticated algorithms for generating, routing, and executing orders
- use of co-location services to minimize network and other types of latencies
- very short timeframes for establishing and liquidating positions
- the submission of numerous orders that are cancelled shortly after submission
- ending the trading day in as close to a flat position as possible

How many stockbrokers does a firm that runs mainly on HFT need to hire? Are there less finance graduates being hired in these firms? Questions that need more data!

3.5.3 The effect on the legal profession

McGinnis and Russell Pearce (2014) have identified five areas of legal services in which algorithms such as machine learning or NLP will greatly disrupt legal services: discovery, legal search, document generation, brief generation, and prediction of case outcomes. Many of these tasks, although cognitively taxing and requiring a substantial education on part of the human labourer, is considered lower level legal work, reserved for interns and recent graduates.

For these fresh graduate lawyers to gain work experience, they start at the bottom. Interns and graduates assist in writing contracts, perform grunt work, generally put their newly minted degrees to use by doing the legwork of the lawyers above them (Alarie et al., 2017). With this, they learn the ins and outs of their trade, and with more knowledge and experience can advance naturally to higher, more complex tasks. However, this work can now be completed by many different algorithms, for example <https://www.luminance.com>, a platform trained by legal experts that can read and understand contracts, read legal documents, and be used for due diligence, compliance, insurance, or contract management. Lumniance themselves argue that this frees lawyers up to focus on what matters. This argument has been posited in many industries facing disruption, for example journalism, and yet you could consider that there are fewer journalists and job opportunities for them. Also, if graduates are no longer receiving hands on experience in the lower levels of being

a lawyer, how will law firms transition to training graduates on higher level matters like handling clients, difficult litigation, and other tasks not currently able to be automated. On-the-job training is crucial to understand any profession.

Another tool is Doxly: www.doxly.com, which is a secure transaction management platform that centralizes checklists, reporting, tracking of documents, tasks, versions, and automates the entire signature management process. They take on many tasks usually performed by transactional attorneys, again perhaps arguing that they free the lawyers to spend more time counselling clients. On one hand, I believe this innovations are vital for increasingly complex fields with huge amounts of data to consider. On the other hand, if graduates are not being hired to take on these tasks and gain necessary work experience, it becomes more difficult for graduates to break into the industry.

Startups such as Doxly or Heretik could hire some graduates, or perhaps inspire entrepreneurial lawyers to create their own, but they serve to substantially shake up the legal profession, potentially harming the economic prospects of many lawyers. They could however, provide advantages to others: superstars in the profession will be more identifiable and will use technology to extend their reach. Also, lawfirms could widen their scope and offer a varied range of services previously unaffordable by harnessing algorithms. These could serve to meet previously unfulfilled legal needs (McGinnis & Pearce, 2013). Another opportunity is that lower-cost, more easily available legal advice, legal help, and perhaps even legal representation to the standard of previously expensive human lawyers could have a revolutionary effect for many previously disadvantaged members of society, who could not afford legal help, who could not understand legal processes, for whom the legal system in different countries and especially when they live in a country of not their birth origin where they are not native speakers.

Notably, lawyers are unlikely, much like doctors, to cease to exist when it comes to the human factor. Algorithms are unlikely to create emotional bonds with clients, and perhaps human counselors are more persuasive against clients that are not making decisions to serve their best interests (McGinnis & Pearce, 2013).

There is even a Norwegian startup in this space: <https://lawtechfactory.com> . They create "automated solutions using state-of-the-art machine learning and natural language processing technology."

3.5.4 But we are lacking data!

Law, medicine, and finance are three fields in which anyone who studies and works hard enough has the potential to improve their standing in society. With the rise of algorithms, this could create further displacement and erode possibilities for social

mobility through job polarization and a lack of opportunities for those entering the three big fields without family contacts.

Although there is a deluge of data in other areas, firm-level data on the adoption of algorithms is lacking. This data would allow researchers to address a host of questions about the extent of adoption of algorithms, under what conditions they are adopted, effects on productivity, what type of firms invest in algorithms, market structure, incentives, and strategies taken by firms that do adopt algorithms (Raj & Seamans, 2018).

Chapter 4

Policy for the uncanny valley

I have noticed that, in climbing toward the goal of making robots appear like a human, our affinity for them increases until we come to a valley, which I call the uncanny valley.

Masahiro Mori

4.1 The state of algorithms, now

Diagnose skin lesions better than dermatologists (Esteva et al., 2017).

Write contracts better, faster, and cheaper than lawyers (Alarie et. al, 2018).

Stocks traded smarter than stockbrokers (Chaboud et al., 2014).

While some fear that robots are coming for physical jobs, they fail to notice that algorithms are already here for the cognitively tasking ones. Task by task. ¹

In this chapter I detail a number of important questions raised by the potential widespread adoption of algorithms, discuss some policy recommendations, and suggest further avenues of future research, ideally performed alongside algorithms.

4.1.1 Who reserves judgement?

There are already algorithms making decisions. Online credit card applications are a good example. If you apply for a credit card with Bank Norwegian, you will get an approval or rejection almost instantaneously. There is not a person sitting there at all hours of the day and night making those judgments. An algorithm makes the decisions. The algorithm exercises judgment on a person's credit score and their income, as reported to the Norwegian tax authority, and based upon set parameters. The algorithm then approves or rejects the application. However, in this case, it is not

¹How do you remove an abortion law? One paragraph at a time.

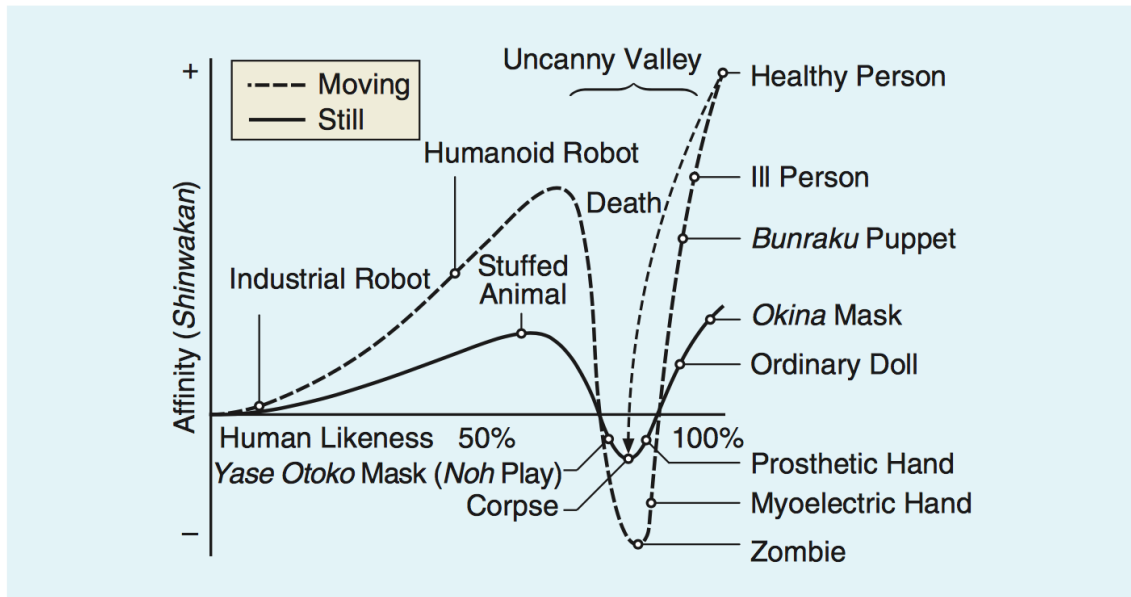


Figure 2. The presence of movement steepens the slopes of the uncanny valley. The arrow's path represents the sudden death of a healthy person. [Translators' note: *Noh* is a traditional Japanese form of musical theater dating to the 14th century in which actors commonly wear masks. The *yase otoko* mask bears the face of an emaciated man and represents a ghost from hell. The *okina* mask represents an old man.]

FIGURE 4.1: From Masahiro Mori's work in 1970 which coined the term the "Uncanny Valley" - *bukimi no tani genshō*

necessarily looking at a judgment, in terms of Agrawal et al. (2017), p.6: "Judgment is the ability to recognize hidden attributes when they arise."

This question is related to the study of Aghion and Tirole (1997) of an investigation of formal versus real authority. In that context, someone who collected the information would report that information to a superior, who may also collect their own information. Part of the role of the boss, or person in charge, is that they cannot pay attention to everything. Hence, the subordinate might be able to exercise real authority over some of the decisions. Moreover, because the subordinate may privately gain from that authority, they would have a greater incentive to collect information. How would algorithms be given authority, and how would we recognize those limitations and plan accordingly?

The ethics of algorithms is a difficult subject. Algorithms have increased autonomy over decisions that have previously been the purview of human beings (e.g., the hiring and firing of employees, whether or not to grant large loan or grant applications, who an autonomous car would kill in a fatal accident). Perhaps some of these situations are misnomers. After all, are you going to judge two people against each other if they make different, subjective decisions? How do you decide which is the preferable subjective decision?

It is difficult to find how much autonomy human beings will retain over decisions made by AI. If a machine learning application that is taught dermatology provides a

different diagnosis than expert human doctors, whom do you consider the expert? Is that a decision that needs to be made by the patient? Is there a third, “blind” expert introduced to act as a judge between the experts and the AI, without knowing who chose what? AI will be trained on human data, in that humans created the contracts that AI learns from, or the laws and legal cases and legal precedents it refers to. Does this mean it will start to write legislation?

What human beings do have control over is in the very inception, in the creation of the original motivations of any kind of artificial type of intelligence. Perhaps even algorithms could regulate their own moral code to a higher degree and a better understanding than human beings. "To the extent that ethics is a cognitive pursuit, a superintelligence could also easily surpass humans in the quality of its moral thinking. However, it would be up to the designers of the superintelligence to specify its original motivations." (Bostrom, 2003, p277). The logic suggests that if we direct the path of superintelligence through its initial, we could have some influence on its path and destination.

Where will the limits of the autonomy of AI lie?

4.1.2 What do we do with the replaced?

What will happen to all of the college graduates when different firms no longer hire people for the grunt work that algorithms can perform for them instead? More broadly, as algorithms begin to diffuse across the economy, it seems likely that a lot of people will get pushed into new positions and many people will be laid off. Just as changing organizational processes takes time, it will also take time to remake the social context in ways that will make it possible to handle these dislocations. Without these kinds of investments (e.g., in education, in relocation assistance), there is a real risk of a public backlash against AI that could dramatically reduce its diffusion rate.

As Rebecca Henderson (2017) notes in her comment on Brynjolfsson et al. (2014), that AI will take time to diffuse because “its adoption will require mastering “adjustment costs, organizational changes and new skills.” Robert Solow’s famous 1987 quote comes into play again, wherein you cannot see change in productivity statistics until firms have actually adopted the changes, created organizations processes around them, and of course most importantly hired human capital that can actually master it. The widespread adoption of algorithms does not only depend on its diffusion but the development of the organizational and human assets to exploit its full potential.

Disruption is an important piece to study, even if 25 years ago Henderson submitted a paper that "suggested that incumbents were fundamentally disadvantaged compared to entrants because they were constrained by old ways of acting and perceiving" and received a rather dismissive letter in return: "Dear Rebecca, you have written a paper suggesting that the moon is made of green cheese, and that economists have too little considered the motions of cheesy planetoids....."(Henderson, 2017, p.2) Although Herbert Simon wrote about organizational change in the 1980s it has taken quite an amount of time to diffuse into economic lexicon, perhaps much like the adoption of algorithms will. After all, academia is not profit-maximizing, so incentives would need to be clear and retraining present to teach academics how to master algorithms.

There needs to be serious study to consider how firms will handle organizational change when it comes to the increasing adoption of algorithms, and how that will affect labour and wages. "As routine-based positions decline due to automation, middle-skilled workers will likely find themselves moving toward less skilled, service-oriented positions, leading to further wage erosion."(Wright Schultz, 2018, p.4) Although economists can model productivity effects, or automation effects, or argue for when wages may rise or fall, serious data is needed to understand the best forms of implementation moving forward. Perhaps this data is best sorted and analyzed by an algorithm coded by motivated economists.

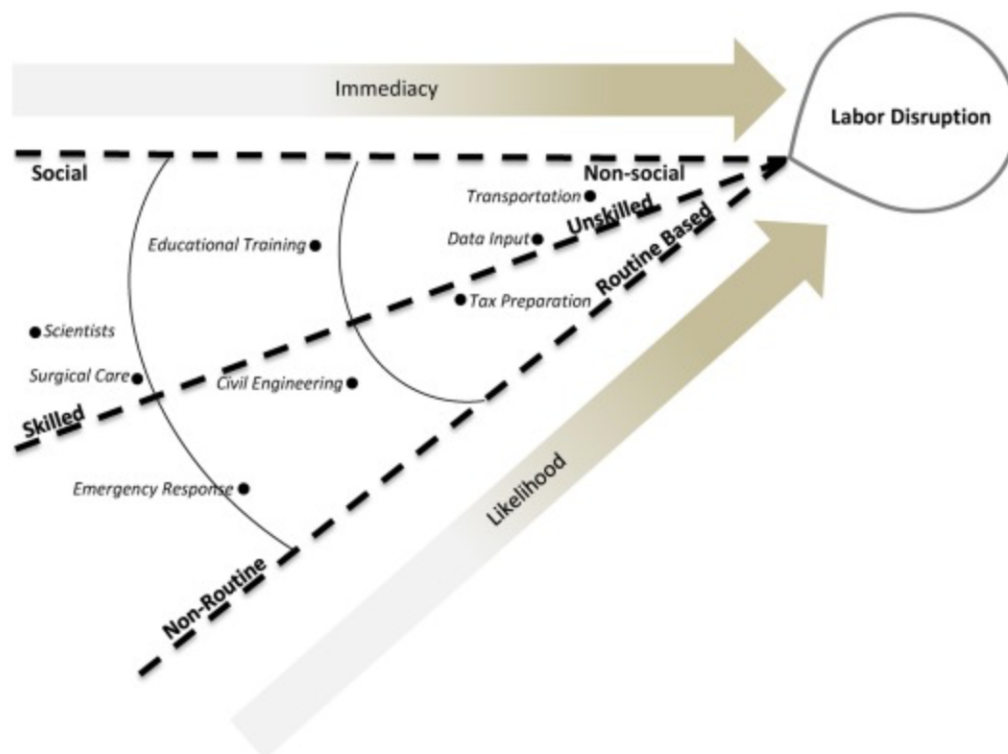


FIGURE 4.2: The tendency towards disruption by algorithms. From Wright & Schultz, 2018

4.1.3 Privacy in the age of AI

Another question of dealing with algorithms is how will a person's privacy be handled when there are fewer and fewer human beings involved in processes? As outlined by Tucker, 2017, there are three ways in which a person's data and their right to privacy is affected.

- **Data Persistence:** What enters the internet may never end up deleted, perhaps even longer than the person that created it.
- **Data Repurposing:** "Once created, such data can be indefinitely repurposed. For example, in a decade's time, parking habits may be part of the data used by health insurance companies to allocate an individual to a risk premium." (Tucker,2017,pp.3)
- **Data Spillovers:** When you submit your DNA to a genealogy database to find out more about your ancestors, about family members, and about your genetic makeup, that data can be used to catch criminals or people who have not consented to having their DNA in a database by virtue of being related to you (GEDMatch, 2018).

Even if you never agree to placing any information on your self on the internet, by virtue of being related to someone who uploads their genetic data to a genealogy database, this may be unavoidable. Persistence and repurposing then raise questions for how to end up avoiding a superintelligence, especially if we keep feeding it data.

4.2 How close are they, really?

The productivity paradox still reigns supreme. That the technology exists does not make it widely available. That it is widely available does not make it successfully implemented. Secondary innovations are arguably as important as the primary ones.

Australia, like Norway, is a small country on the global stage, and like Norway considers the future through the lens of setting examples but not necessarily steering the future. In the megatrend report by Reeson et al. (2016) on "Tomorrow's digitally enabled workforce", the authors note that although Australia's workforce is continually changing and evolving, this current period in history is different. There are a combination of factors at play that create a potential 'perfect storm' for massive and speedy transitions than previously experienced.

The full impact of exponential and/or steep growth in computing power, device connectivity, data volumes and artificial intelligence is yet to be felt within Australia's labour market. Many reports, from the OECD (2016) to Frey and Osborne (2013) imply that there are a large cohort of existing jobs that are likely to be automated which means many new jobs need to be created to take on the displaced workforce.

Perhaps a solution that the market itself offers is peer-to-peer (P2P) marketplace (Upwork, Freelancer, Kaggle, Task, Rabbit et al.), where freelancers or skilled individuals can find short-term contracts or projects to work on (Reeson et al., 2016). Another point to note is that we live in an increasingly globalised labour market where individuals have the opportunity to work remotely. For example, a program called Tulsa Remote, sponsored by the George Kaiser Family Foundation, offers 10 000 USD upfront to skilled individuals to move to the city of Tulsa and bring with them their remote work (Tulsa Remote, 2018). The foundation recognizes an opportunity in reviving the city of Tulsa without necessarily creating new job opportunities there. Perhaps remote villages and towns in Norway could take inspiration.

Much like Norway, Australia is entering a post-mining boom era of its economic development and so needs to be positioned for diversification into services, knowledge, and innovation exports. Norway, like Australia, is a small nation that can be well positioned to pivot to new opportunities. As relatively small economies, Australia and Norway will both likely be a net importers of technology outside of their specific realms of expertise (for example off-shore drilling platforms in Norway). For knowledge and technology, large global platforms that have scaled offer far greater affordability, and can be adapted to local conditions.

Workforce transitions, in how individuals move from one job to another and how industries move from one labour market structure to another, are crucial to understand and plan for. Even if change is inevitable, the future of work and working with algorithms is not. For smaller countries with highly educated workforces, a narrative of the future can be constructed that empowers individuals, communities, companies and governments to better identify and implement transition pathways that achieve good outcomes. Job changes and multiple careers in a working life are far more likely to become the norm; this pattern has been happening for a while.

Workers will need the capability to handle a career dead-end (or job loss) and create their own job in another space. This is a truth facing even academics, or graduates with PhDs no longer finding professor-track positions. To be a resilient individual, you need to develop character traits and mindsets which are learned over time and potentially through difficult experiences. These are not necessarily easily built into a structured curriculum. Soft skills may be increasingly vital for people to thrive in tomorrow's labour market (Stevenson, 2018).

At a panel of economists questioned about the future of work, one economist noted that those not in the labour force are often unhappy and have an inclination to take opioids (Stevenson, 2018). Naturally, this perhaps is true for particular cohorts, and looking at feminist economics it may not be the only way to describe society. I agree though, that the notion of work as purpose or meaning in many people's lives (Morse Weiss, 1955) is crucial in understanding the future of how economies

will adopt algorithms, and the potential effects it may have on societies around the world.

4.2.1 What will the use us for?

That algorithms can take on what we consider to be cognitively taxing tasks, does not mean that there are tasks beyond them (Brynjolfsson & Mitchell, 2017). Three areas in which algorithms are lacking prowess, as outlined in the OECD Report on Automation, Skills Use, and Training (2018). The first is social intelligence, the ability to effectively negotiate complex social relationships, whether that means caring for others, recognising cultural sensitivities, understanding power dynamics, and other interpersonal aspects of working life. Naturally, the question arises of the algorithmic equivalent: how well do algorithms “talk” to one another? How well can they be integrated? Will an algorithm specialize in integrations between the multiple tasks that create a job?

Another field that human beings will take the lead is creativity. Cognitive intelligence, any type of considerable generalization, and application outside of strictly narrow fields (unless by perhaps an integration between millions of narrow algorithms), is lacking in today’s abilities algorithms and artificial intelligence. Algorithms can complete one task well, often better than their human counterparts, with enough training and enough good data (Brynjolfsson & Mitchell, 2017). However, creativity involves creating new connections, seeking new patterns, copying and melding ideas into a brand new solution, often never before imagined or understood. This is something human beings excel at. Complex reasoning is another questionable concept for algorithms.

Finally, although Boston Dynamics has created many an amazing performance, perception and manipulation is still out of reach when it involves physical tasks carried out in an unstructured work environment (Murray, 2017). Robots would need situational awareness, and if working alongside humans, efficient use of algorithms such as NLP to follow and understand human direction (Duvallet, 2015). Robots can backflip, or walk down stairs, but stacking boxes onto a palette, and loading those onto a truck is a mammoth task for an algorithm/robot combination²(Nelson et al., 2019). The implementation of robots into the workforce is a topic outside the scope of this investigation, but it bears mentioning.

There are many opportunities to work with algorithms, many professions would benefit from people having more face to face time with individuals that they work with or for, rather than with the documents or tasks around those individuals. However, this is a philosophical shift for many professions to take, and difficult to see if

²I have avoided talking about robots as much as possible in this work as I have wanted to focus on the effect that algorithms themselves will have on the future of labour without falling into the dystopian traps of science fiction.

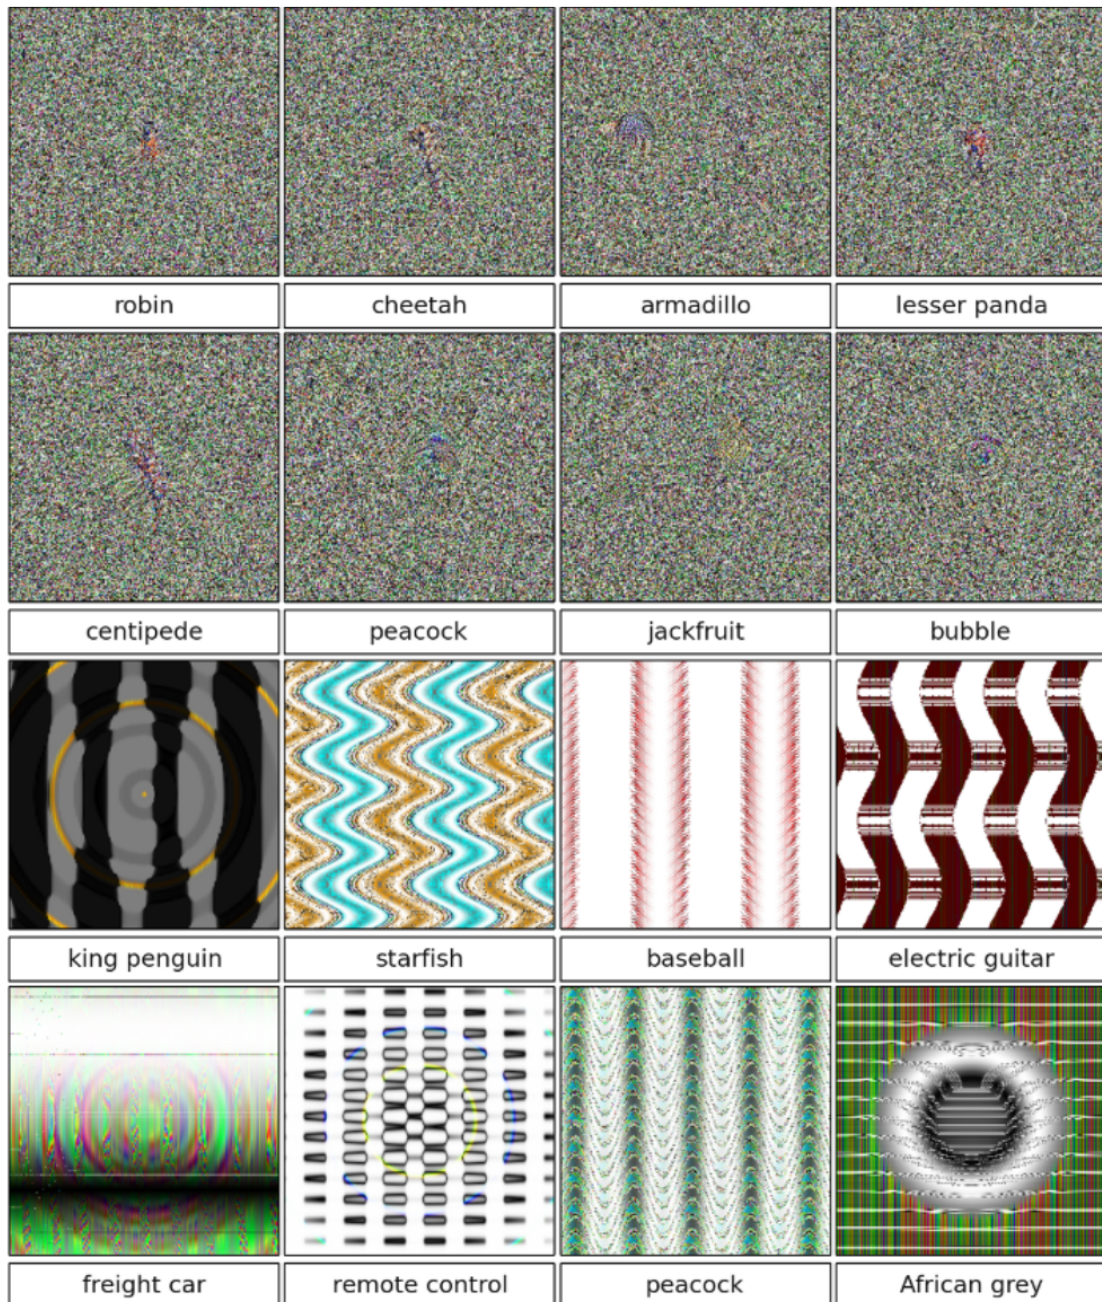


Figure 1. Evolved images that are unrecognizable to humans, but that state-of-the-art DNNs trained on ImageNet believe with $\geq 99.6\%$ certainty to be a familiar object. This result highlights differences between how DNNs and humans recognize objects. Images are either directly (*top*) or indirectly (*bottom*) encoded.

it will be in the interest of profit-maximizing firms to do so, if the easier and cheaper option is merely to replace workers.

4.2.2 Opportunity and risk are indeed twins: with new opportunities come new problems.

There was an incident on the U.S. stock market on May 6th, 2010 which has since been named the "Flash Crash" where, in the course of around 30 minutes, U.S. stock market indices, stock index futures, options, and exchange-traded funds experienced a sudden price drop of more than 5 percent, followed by a rapid rebound. It was a brief period of extreme volatility (Kirilenko et al., 2017). Incidents like this highlight that a world increasingly run with and by algorithms could have repercussions when it comes to increased volatility in stock markets, which play a role in political choices that individuals make.

Although algorithms bring new opportunities, like algorithmic trading or better understanding of the masses of data that are created each day, by allowing algorithms increased decision and judgement power, there is an increased risk of volatility or even injury. If the robot surgeon hand makes a mistake, who is liable? If the algorithm misses a clause in a contract, will the intern be trained well enough to pick it up? It is perhaps a little too easy to argue for the irrelevance of the human labourer in a world of algorithms that can end up with serious biases (Armantier & Copeland, 2012).

4.3 Regulation

Multi-national corporations have made an artform of successfully playing different governments against each other to achieve their best terms, whether it comes to taxation, law, or regulation. With algorithms taking over more and more current tasks, how this will play out has a playbook. What is the dream of a corporation, if not to have no staff? No labour costs? No pesky unions, or safe working laws, or employee satisfaction, or holidays, or payroll, indeed there is much to be free of with a skeleton staff. With the exponential growth of algorithms, the tasks they can complete, and the breadth of application, this is sooner reality than science fiction. Regulation, and good regulation that understands the risks and benefits of algorithms as applied in different professions or to different tasks will be perhaps the only way to manage this once the floodgates well and truly open.

However, if you are to regulate all of these narrow intelligences, you need to understand it. And what better way to understand it than with an algorithm designed for doing so. What an algorithm will soon do better than researchers, than academics, than government officials is to find and process information. Naturally there are

people at one end asking questions, or seeking information, but even with enough minds to process all the complexities that arise from deep interference of algorithms, this is also a matter of time. In order to understand the advances of algorithms into different tasks, an algorithm will need to do that research. It will need to process many different kinds of information to show regulators what areas are changing quickly and need attention in order for the transition, if there is one, to artificial intelligence in different professions is as smooth as possible. It will be necessary in the effective regulation of algorithms.

There has been little technology that has been so broad in scope as algorithms. To regulate algorithms, to understand all the different and perhaps insidious influences or changes that come, even corporations such as Facebook, Google, Microsoft, and IBM work through Partnership on Artificial Intelligence to Benefit People and Society (Wright & Schultz, 2018). This partnership aims to address ethical issues and improve society's understanding of artificial intelligence. This is not a zero sum game, AI used to process information will be incredibly beneficial and productive for many professions that are built up of tasks involving research and calculations. With the correct checks and balances, government committees, task forces, and in the future certainly departments and ministerial positions will be important in a cohesive transition to a society where AI and algorithms play a vital role in our daily lives.

4.4 The control problem

Writing about algorithms and about artificial intelligence lends itself to big, often scary ideas lurking in the corners. One of these ideas is the control problem. How do you build a superintelligence that does not harm the project's interests (Bostrom, 2014)? The control problem arises when there is no way for a human to insure against existential risks before an artificial general intelligence AGI becomes superintelligent, either by controlling its capabilities or its motivations (Gans, 2018).

One classic example was presented by Bostrom (2003), called the paperclip problem. Even when an AGI has a goal that is arbitrary, such as manufacturing as many paperclips as possible, if an AGI is superintelligent it can quickly spiral out of control, for example transforming the earth and space into paperclip manufacturing facilities. That the AGI did not start with a motive anywhere near domination does not mean it would not result in monopolizing all available resources. The idea goes as dark as to say that this could even result in the subjugation or elimination of humans by the paperclip making AGI to achieve its goal. Bostrom (2014) presents that AI researchers are yet to sufficiently come up with methods of control that would ensure such outcomes did not take place.

The control problem is a superintelligence version of the classic principal-agent problem whereby a principal faces decisions that seek to ensure an agent with different goals acts in the interest of the principal (Gans, 2018). Although this problem is familiar to economists, theory is imperfect and ill-equipped to resolve the risks associated with control of a superintelligence. For those most fearful of this, they say that the only tool that can prevent harm by a powerful agent is the removal of their agency. To remove the agency of a superintelligence would be to not create it at all (Barrat, 2013).

It is a strange sort of argument indeed. But a novel way to handle the control problem of artificial intelligence. There are many lenses and frames through which you can understand algorithms and AI, and the control problem is a good model to understand the consequences and complexity that arise from having an ultimate invention that could improve itself recursively ad infinitum (Good, 1966).

If there is a machine, algorithm, or superintelligence that could self-improve, perhaps the only limitation we could impose is a lack of resources available to it. The basic idea is that machines will be able to self-improve by building even better machines or as in the case of the paperclip problem, a machine that works out ways of appropriating more resources. In AI research this goes under the term recursive self-improvement (Yudkowsky, 2007). The central argument of Gans (2018) is that limiting AI, for example by not allowing it property rights, would protect AI in the sense that human beings would allow it to exist. Looking at the progress of narrow intelligence across so many different fields, this author feels its honestly only a matter of time and connections between algorithms before a superintelligence is created.

4.5 Global reactions

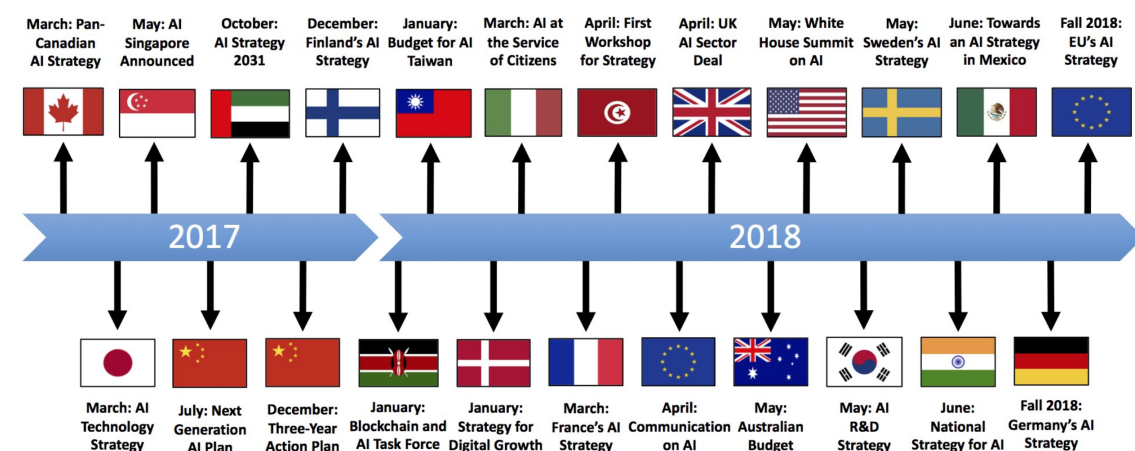


FIGURE 4.4: Global responses to AI. Dutton, (2018)

Across the world, different governments are reacting to changes in their workforces, in global competition, and in attempting to understand the future of automation and of work. For example, in China, Baidu (China's Google), Tencent (China's Facebook on steroids), and the Chinese government have created the Chinese Panopticon (Tan & Ding, 2018). The Chinese government seem well aware of the fact that AI technologies have the potential to further entrench the near-monopoly power of technology giants.

Apps such as Tencent's Wechat, what Facebook aspires to be, can collect data on user searches, navigation patterns, reading interests, and payments in a single platform. This "superapp" could swallow up products that provide even the most specialized AI products and services (Tan & Ding, 2018). Nations such as China could be more willing to designate certain large companies as national champions as it is potentially easier to manage a smaller number of AI giants than a higher number of specialized firms (Dutton, 2018).

Moving across the Pacific Ocean, Canada was the first nation to create an AI strategy, with the Pan-Canadian AI strategy over five years as outlined in the 2017 federal budget. Led by CIFAR, the 125 million dollar plan outlines different areas to focus on: increased education, create interconnected nodes of scientific excellence, develop global thought leadership, and support a national research community on artificial intelligence (CIFAR, 2018).

Leaping over Canada across the Atlantic, we come to the EU, where the aim is to increase the EU's technological and industrial capacity while harnessing and encouraging AI uptake by the public and private sectors. Another key initiative as outlined in the Communication on Artificial Intelligence for Europe (2017) is to prepare the citizens of the EU for the socioeconomic changes brought about by AI, many of which I have outlined earlier in this work. Finally, it is key to ensure that an appropriate ethical and legal framework is in place. Vital initiatives include a commitment to increase the EU's investment in AI from 500 million in 2017 to 1.5 billion by the end of 2020, and the creation of the European AI Alliance (AI for Europe, 2017).

In the USA, multinational corporations rule. All of Silicon Valley is in on the algorithms game. The government of the United States recently announced a Select Committee on Artificial Intelligence which will operate through the National Science and Technology Council (White House, 2017). They see the potential of harnessing AI: "Artificial intelligence holds tremendous potential as a tool to empower the American worker, drive growth in American industry, and improve the lives of the American people. Our free market approach to scientific discovery harnesses the combined strengths of government, industry, and academia, and uniquely positions us to leverage this technology for the betterment of our great Nation." Michael Kratsios, Deputy Assistant to the President for Technology Policy. (White House, 2018).

The purpose of the committee seems to be to understand how AI will affect American society from an employment, economic, national security, and health perspective. Understanding AI is an important and vital step for all governments around the world to take, and although the U.S. specifically states that “Overly burdensome regulations do not stop innovation—they just move it overseas.” there is certainly regulation that will need to be enforced to best understand and harness AI while protecting workers from excessive displacement.

Countries in Scandinavia must seriously consider how they treat algorithms. The decisions on the vast changes coming to many people in the workforce need to be planned today. Can a government own algorithms? Can they outsource parts of them to the private sector? What will the public sector look like, run by algorithms?

In Australia, the greatest problem will be internet speed. There cannot be a future of algorithms and cloud computing without wildly faster internet. Perhaps algorithms can be created that adapt to the conditions of slower, intermittent, or unreliable internet. This may be an entirely new industry that forms.

If countries with slower internet, and perhaps no realistic nor cost effective to quickly create faster internet, want to use algorithms, there will be private actors that solve this equation in any number of ways. For example, Babylon, the UK health app that allows patients in the NHS to videocall doctors, has an agreement with the Government of Rwanda to assist in providing healthcare. As smartphones are not widely available in the nation of Rwanda, patients call nurses who work with the chatbot, asking questions as prompted by the screen and interpreting answers given by patients in a way that the app understands (Babylon, 2018). Initiatives such as this may help thwart the productivity effect as Zeira, and later Açemoğlu & Restrepo outline of increased capital availability as a means of widening between nation differences (Zeira, 1998).

4.6 What's next?

There are many questions that this thesis only touches on. These questions may need to be explored in more depth. Other questions are already being explored in the fields of AI. These questions include the future of work and the future of society.

It is recommended that economists consider policy advice for governments on the topics of AI, algorithms, and the replacement of labour. Although this work is optimistic in how the role of work will change, the quantity of work available for a growing global population is a very pertinent question.

Robots were not covered in the scope of this thesis. However, there are many questions around how robotic labour will be used in the future. There are many difficult, menial tasks, often dangerous tasks, that are well suited to robotic work. These tasks

are often comprised of many low-paid, labour-intensive jobs, which are a broad part of the employment opportunities for many people around the world without an extensive formal education. Without access to higher education or opportunities related to higher education, what type of work will these people have to do in the future?

The discussions around a Universal Basic Income also needs to assess situations and relevance to the societies of people who are highly educated, but lack opportunities (Laukkanen et al., 2018). This lack of opportunities may be due to their local economy not offering these opportunities or because many of the tasks that comprise their job role have been replaced by algorithms, AI, robots, or a combination. Thus, there is less demand for their skilled labour. More research is needed in the areas of re-skilling, re-training, career changes, and lifelong education. These studies are important to determine how future high-level economies function. Algorithms will cover a lot of the grunt-work type tasks, while humans will conduct the creative tasks. With many new fields opening up, and the accelerated progress of economies and societies, humans will need frequent education and up-skilling, a substantial increase in the soft skill of career resilience (Rossier et al, 2017). It might be that taking ongoing higher education is an essential task in a corporate setting.

Governments need to consider the ownership of algorithms, AI, and robots. Algorithms, AI and robots may be labour, or labour-replacing. If this power is in the hands of multinational corporations, there may be little interest in a nations' wellbeing (provided you can make the assumption that a government has an interest in a nations' wellbeing, especially for the majority of its citizens). This may result in an acceleration of income and social inequality if algorithms and the subsequent wealth they generate is held in by a shrinking group of technocrats.

If people are to work alongside algorithms, the public sector is an excellent place for that to occur. Naturally, there are always incentive issues with public funding in that corporations often deliver better products with higher incentives.

If the governments major income source is taxation, then governments naturally have a high incentive to keep people working. However, they need to balance taxation with work coming from sources outside of the government, unless the product the people employed by the government.

4.6.1 We've been at crossroads before

Massive change can occur within generations. With the second world war, women were "freed" to work. This change led to a massive demographic shift, where a majority of women work outside of their households (Stevenson, 2008). Technological change occurred in a way that crowded out homemade goods and crowded in women's labor force participation (Stevenson Wolfers, 2007). Professions have also

changed. New job titles abound, and many more aspects of work are now measured (Stevenson, 2008). The industrial revolution, then digitalization, and now the proliferation of algorithms. An examined life is worth living, and so an examined future is worth having.

Chapter 5

Conclusion

In this investigation I have discussed at length the effect that algorithms are having on the workforce through different empirical sources. The future is always closer than we think, and understanding how algorithm will affect the workforce is presented through a discussion of a model of task-based automation. Although this future is not inevitably bad, without good planning, research, data, opportunities, and local considerations, there is a high risk of considerable worker displacement and replacement. This is especially if individuals are not retrained to work with algorithms. Career resilience will be vital for the near future.

We will work with algorithms. Not just for, not instead of, not against, but with them. How algorithms differ from perhaps the personal computer or other technological inventions which have augmented human labour is that algorithms are not limited to being just a tool. They will be more like an intelligent side-kick, or a clever chatbot, or an analysis program that creates information for human beings to create with, to analyze, to process, and to be wildly more productive with, given the correct training and testing of the algorithms.

Appendix A

References

A.1 Papers

Acemoglu, D., Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In Handbook of labor economics (Vol. 4, pp. 1043-1171). Elsevier.

Acemoglu, D., Restrepo, P. (2018). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review*, 108(6), 1488-1542.

Acemoglu, D., Restrepo, P. (2018). Artificial Intelligence, Automation and Work (No. w24196). National Bureau of Economic Research.

Agrawal, A. K., Gans, J. S., Goldfarb, A. (2018). Exploring the impact of artificial intelligence: Prediction versus judgment (No. w24626). National Bureau of Economic Research.

Alarie, B., Niblett, A., Yoon, A. H. (2018). How artificial intelligence will affect the practice of law. *University of Toronto Law Journal*, 68(supplement 1), 106-124.

Armantier, O., Copeland, A. M. (2012). Assessing the Quality of 'Furfine-Based' Algorithms.

Arntz, M., Gregory, T., Zierahn, U. (2016). The risk of automation for jobs in OECD countries.

Athey, S. (2018). The impact of machine learning on economics. In *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Autor, D., Salomons, A. (2018). Is automation labor-displacing? Productivity growth, employment, and the labor share. *Brookings Papers on Economic Activity*.

Bresnahan, Timothy F., and M. Trajtenberg (1995), "General Purpose Technologies 'Engines of Growth'?" *Journal of Econometrics* 65 (1):83-108.

Brynjolfsson, E. and Mitchell, T. (2017), "What can machine learning do? Workforce implications.", *Science*, 358(6370), pp.1530-1534.

- Brynjolfsson, E., Rock, D., Syverson, C. (2017). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In *Economics of Artificial Intelligence*. University of Chicago Press.
- C. B. Frey and M. A. Osborne, *The future of employment: How susceptible are jobs to computerisation?* Oxford, United Kingdom: Department of Engineering Science, University of Oxford, 2013.
- Chaboud, A. P., Chiquoine, B., Hjalmarsson, E., Vega, C. (2014). Rise of the machines: Algorithmic trading in the foreign exchange market. *The Journal of Finance*, 69(5), 2045-2084.
- Chambers, J. D., Mingay, G. E. (1966). *Agricultural revolution 1750-1880*.
- Duvallet, F. (2015). *Natural Language Direction Following for Robots in Unstructured Unknown Environments* (No. CMU-RI-TR-15-01). CARNEGIE-MELLON UNIV PITTSBURGH PA ROBOTICS INST.
- Erev, I., Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, 848-881.
- Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115.
- Feinstein, C. H. (1998). Pessimism perpetuated: real wages and the standard of living in Britain during and after the industrial revolution. *The Journal of Economic History*, 58(3), 625-658.
- Feldman, R., Dagan, I. (1995, August). Knowledge Discovery in Textual Databases (KDT). In *KDD* (Vol. 95, pp. 112-117).
- Foucault, T., Hombert, J., Roşu, I. (2016). News trading and speed. *The Journal of Finance*, 71(1), 335-382.
- Gans, J. S. (2018). *Self-Regulating Artificial General Intelligence* (No. w24352). National Bureau of Economic Research.
- Grace, Katja (2013), "Algorithmic Progress in Six Domains", Technical report 2013-3. Berkeley, CA: Machine Intelligence Research Institute.
- Gregory, T., Salomons, A., Zierahn, U. (2016). Racing with or against the machine? Evidence from Europe, *Utrecht School of Economics Discussion Paper Series* 16-05
- Goos, M., Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in Britain. *The review of economics and statistics*, 89(1), 118-133.
- Guerreiro, J., Rebelo, S., Teles, P. (2017). Should Robots be Taxed? (No. w23806). National Bureau of Economic Research.

Henderson, R. (2018). Comment on 'Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics'. In *Economics of Artificial Intelligence*. University of Chicago Press.

Hobsbawm, E. J. 1952. *The Machine Breakers*. Past and Present 1 (February): 57–70.

Jarrow, R. A., Protter, P. (2012). A dysfunctional role of high frequency trading in electronic markets. *International Journal of Theoretical and Applied Finance*, 15(03), 1250022.

Jones, C. I., Romer, P. M. (2010). The new Kaldor facts: ideas, institutions, population, and human capital. *American Economic Journal: Macroeconomics*, 2(1), 224-45.

Keynes, John Maynard (1930), "Economic Possibilities for our Grandchildren," in *Essays in Persuasion*, Harcourt Brace, pp. 358-373.

Kozhan, R., Tham, W. W. (2012). Execution risk in high-frequency arbitrage. *Management Science*, 58(11), 2131-2149.

Kirilenko, A., Kyle, A. S., Samadi, M., Tuzun, T. (2017). The Flash Crash: High-frequency trading in an electronic market. *The Journal of Finance*, 72(3), 967-998.

Kriegeskorte, N. (2015). Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annual review of vision science*, 1, 417-446.

Korinek, Anton and Joseph Stiglitz (2017), "Artificial Intelligence and Its Implications for Income Distribution and Unemployment," NBER Working Paper w24174.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40.

Laukkanen, E. (2018). Andy Stern (Author) Lee Kravitz (Contributor): *Raising the Floor: How a Universal Basic Income Can Renew Our Economy and Rebuild the American Dream*. Basic Income Studies.

LeCun, Yann, Bengio, Yoshua, Hinton, Geoffrey. (2015). Deep learning. *nature*, 521(7553), 436.

Lemley, J., Bazrafkan, S., Corcoran, P. (2017). Deep Learning for Consumer Devices and Services: Pushing the limits for machine learning, artificial intelligence, and computer vision. *IEEE Consumer Electronics Magazine*, 6(2), 48-56.

Leontief, W. (1983) National perspective: The definition of problem and opportunity, in: *National Academies, The Long-term Impact of Technology on Employment and Unemployment: A National Academy of Engineering Symposium*

Liddy, E. D. (2001). *Natural language processing*.

- Llewellyn Evans, G. (2017). Disruptive technology and the board: The tip of the iceberg. *Economics and Business Review*, 3(1).
- Mann, K., Püttmann, L. (2017). Benign Effects of Automation: New Evidence From Patent Texts.
- Manyika, J., Chui, M., Miremadi, M., Bughin, J., George, K., Willmott, P., Dewhurst, M. (2017). *A Future that Works: Automation, Employment, and Productivity*. McKinsey Global Institute.
- McAfee, A., Brynjolfsson, E. (2016). Human work in the robotic future: Policy for the age of automation. *Foreign Aff.*, 95, 139.
- McGinnis, J. O., Pearce, R. G. (2013). The great disruption: How machine intelligence will transform the role of lawyers in the delivery of legal services. *Fordham L. Rev.*, 82, 3041.
- Menkveld, A. J. (2013). High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4), 712-740.
- Mokyr, J., Savin, N. E. (1976). Stagflation in historical perspective: the Napoleonic wars revisited. *Research in Economic History*, 1, 198-259.
- Mori, M. (1970). The uncanny valley. *Energy*, 7(4), 33-35.
- Morse, N. C., Weiss, R. S. (1955). The function and meaning of work and the job. *American Sociological Review*, 20(2), 191-198.
- Murray, R. M. (2017). *A mathematical introduction to robotic manipulation*. CRC press.
- Nelson, G., Saunders, A., Playter, R. (2019). The PETMAN and Atlas Robots at Boston Dynamics. *Humanoid Robotics: A Reference*, 169-186.
- Nguyen, A., Yosinski, J., Clune, J. (2015). Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 427-436).
- Nordhaus, W. D. (2015). Are we approaching an economic singularity? Information technology and the future of economic growth (No. w21547). National Bureau of Economic Research.
- O'Rourke, K. H., Rahman, A. S., Taylor, A. M. (2013). Luddites, the industrial revolution, and the demographic transition. *Journal of Economic Growth*, 18(4), 373-409.
- Pajarinen, M., Rouvinen, P. (2014). Computerization threatens one third of Finnish employment. *Etla Brief*, 22(13.1), 2014.
- Raj, M., Seamans, R. (2018). *AI, Labor, Productivity, and the Need for Firm-Level Data*. In *Economics of Artificial Intelligence*. University of Chicago Press.

- Ranjan, R., Sankaranarayanan, S., Bansal, A., Bodla, N., Chen, J. C., Patel, V. M., ... Chellappa, R. (2018). Deep learning for understanding faces: Machines may be just as good, or better, than humans. *IEEE Signal Processing Magazine*, 35(1), 66-83.
- Robinson, J. (1953). The production function and the theory of capital. *The Review of Economic Studies*, 21(2), 81-106.
- Rossier, J., Ginevra, M. C., Bollmann, G., Nota, L. (2017). The importance of career adaptability, career resilience, and employability in designing a successful life. In *Psychology of Career Adaptability, Employability and Resilience* (pp. 65-82). Springer, Cham.
- Richardson, G. (2008). Brand names before the industrial revolution (No. w13930). National Bureau of Economic Research.
- Stevenson, B. (2018). AI, Income, Employment, and Meaning. In *Economics of Artificial Intelligence*. University of Chicago Press.
- Stevenson, B., Wolfers, J. (2007). Marriage and divorce: Changes and their driving forces. *Journal of Economic perspectives*, 21(2), 27-52.
- Tan, J. Z., Ding, J. (2018). AI governance through AI markets.
- Tucker, C. (2017). Privacy, Algorithms and Artificial Intelligence. In *Economics of Artificial Intelligence*. University of Chicago Press.
- Turing, A. (1950). Computing Machinery and Intelligence. *Mind*, 59(236), 433-460.
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2), 3-28.
- Wright, S. A., Schultz, A. E. (2018). The rising tide of artificial intelligence and business automation: Developing an ethical framework. *Business Horizons*.
- Yongfeng Qian, Jiayi Lu, Yiming Miao, Wen Ji, Renchao Jin, Enmin Song, "AIEM: AI-enabled affective experience management", *Future Generation Computer Systems* 89 (2018) 438-445
- Yudkowsky, E. (2007). Levels of organization in general intelligence. In *Artificial General Intelligence*, ed. B. Goertzel C. Pennachin, 389-501.
- Zeira, J. (1998). Workers, machines, and economic growth. *The Quarterly Journal of Economics*, 113(4), 1091-1117.
- Zenati, M. A. (2001). Robotic heart surgery. *Cardiology in review*, 9(5), 287-294.
- Zheng, S., Song, Y., Leung, T., Goodfellow, I. (2016). Improving the robustness of deep neural networks via stability training. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4480-4488).

A.2 Books

Barrat, James. (2013). *Our Final Invention: Artificial Intelligence and the End of the Human Era*. New York: St. Martin's Press.

Bostrom, N. (2003). Ethical issues in advanced artificial intelligence. *Science Fiction and Philosophy: From Time Travel to Superintelligence*, 277-284.

Gersting, J., Schneider, G. M. (1995). *An Invitation to Computer Science*.

Goertzel, B. (2007). *Artificial general intelligence (Vol. 2)*. C. Pennachin (Ed.). New York: Springer.

Kahneman, (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.

Siegel, M. (2018). Text Mining in Economics. In *Semantic Applications* (pp. 63-73). Springer Vieweg, Berlin, Heidelberg.

A.3 Reports

Blit, J., Amand, S. S., Wajda, J. (2018). *Automation and the Future of Work: Scenarios and Policy Options*.

Bureau of Labor Statistics (2017): Industry Output and Employment Projections https://www.bls.gov//www.bls.gov/emp/ep_data_occupational_data.htm

COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT, THE EUROPEAN COUNCIL, THE COUNCIL, THE EUROPEAN ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS (2017), <https://ec.europa.eu/digital-single-market/en/news/communication-artificial-intelligence-europe>, accessed 20th September 2018.

Domo Report (2018), *Data Never Sleeps 5.0*, <https://www.domo.com/learn/data-never-sleeps-5?aid=ogsm0725171sf100871281=1>, accessed 28th November 2018.

Dutton, P. (2018), *An Overview of National AI Strategies Politics + AI*, Doc: <https://medium.com/policy-ai/an-overview-of-national-ai-strategies-2a70ec6edfd>

How we use data and analytics (2018), <https://www.ato.gov.au/about-ato/managing-the-tax-and-super-system/insight-building-trust-and-confidence/how-we-use-data-and-analytics/>, accessed 3rd July 2018.

McKinsey Global Institute (2017), "What the future of work will mean for jobs, skills, and wages," November: <https://www.mckinsey.com/global-themes/future-of-organizations-and-work/what-the-future-of-work-will-mean-for-jobs-skills-and-wages>

Oxford Martin School (2016): "Technology at Work v2: The Future Is Not What It Used To Be", available at (<http://www.oxfordmartin.ox.ac.uk/downloads/reports/CitiGPTechnologyWork>)

SEC (2010). Concept release on equity market structure. Release No. 34-61358; File No. S7-02-10.

Stone, p., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., Press, W., Saxenian, A., Shah, J., Tambe, M., and Teller., A. (2015) "Artificial Intelligence and Life in 2030." One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, Stanford University, Stanford, CA, September 2016. Doc: <http://ai100.stanford.edu/2016-report>. Accessed: October 28, 2018.

Tulsa Remote (2018). <https://tulsaremote.com/hero>. Accessed: November 28, 2018.

Uliyar, S. (2017), A Primer: Oracle Intelligent Bots - Powered by artificial intelligence, <http://www.oracle.com/us/technologies/mobile/chatbots-primer-3899595.pdf>

White House Hosts Summit on Artificial Intelligence for American Industry, <https://www.whitehouse.gov/white-house-hosts-summit-artificial-intelligence-american-industry/>, accessed 5th July 2018.

World Economic Forum (2016), "The Future of Jobs," Davos: Global Challenge Insight Report: <http://reports.weforum.org/future-of-jobs-2016/>

Appendix B

Derivations

B.0.1 Derivations for the Basic Model

Let us build the model from its lego block pieces. We start with price, $p(x)$, the price of task x :

$$p(x) = \begin{cases} \frac{R}{\gamma_A(x)} & \text{if } x \in [0, I] \\ \frac{W_Z}{\gamma_Z(x)} & \text{if } x \in (I, F) \\ \frac{W_H}{\gamma_H(x)} & \text{if } x \in [F, N] \end{cases} \quad (\text{B.1})$$

Then we have demand for task x as given by:

$$y(x) = \frac{Y}{p(x)} \quad (\text{B.2})$$

Then there are three types of demand. The demand for algorithms in task x is:

$$a(x) = \begin{cases} \frac{Y}{R} & \text{if } x \in [0, I] \\ 0 & \text{if } x \in (I, F) \\ 0 & \text{if } x \in [F, N] \end{cases} \quad (\text{B.3})$$

The demand for the combination of humans and algorithms in task x is:

$$z(x) = \begin{cases} 0 & \text{if } x \in [0, I] \\ \frac{Y}{W_Z} & \text{if } x \in (I, F) \\ 0 & \text{if } x \in [F, N] \end{cases} \quad (\text{B.4})$$

and the demand for only labour in task x is:

$$h(x) = \begin{cases} 0 & \text{if } x \in [0, I] \\ 0 & \text{if } x \in (I, F) \\ \frac{Y}{W_H} & \text{if } x \in [F, N] \end{cases} \quad (\text{B.5})$$

On the flipside of demand we have supply. We then aggregate each factor and make it equal to its respective supply.

The supply of algorithms:

$$A = \frac{Y}{R}(I - N + 1) \quad (\text{B.6})$$

The supply of humans who work with algorithms, and their algorithms:

$$Z = \frac{Y}{W_Z}(F - I) \quad (\text{B.7})$$

The supply of humans who don't work with algorithms:

$$L - Z = \frac{Y}{W_H}(N - F) \quad (\text{B.8})$$

The expression for total labour is then made up of $L - Z + Z = L$.

I define three task types to reflect the changing economy of the future. Labour, in today's economy, interacts with machines to complete the allocated tasks; indeed, some tasks can only be completed by combining humans and computers (e.g., sending an email, taking a meeting using video-calling technology, performing a mammogram), let alone any necessary use of the internet for communications or information.

I assume the supply of labour, L , and algorithms, A , are fixed and inelastic. This allows for changes in demand to affect not the level of employment but the share of labour in the wage and national income. Later I discuss algorithms as capital or labour or neither, however here in the model they play the function of capital.

Now let's look at deriving total output. I have normalized the price of the final good to 1 as a numeraire, following Açemoğlu & Restrepo (2018).

$$\int_{N-1}^N \ln p(x) dx = 0 \quad (\text{B.9})$$

Substituting in the expressions for $p(x)$ the equation becomes:

$$\int_{N-1}^I [\ln R - \ln \gamma_A(x)] dx + \int_I^F [\ln W_Z - \ln \gamma_Z(x)] dx + \int_F^N [\ln W_H - \ln \gamma_H(x)] dx = 0 \quad (\text{B.10})$$

Using the equations derived for A , Z , and $L - Z$:

$$\begin{aligned} & \int_{N-1}^I [\ln Y - \ln(A/(I - N + 1)) - \ln \gamma_A(x)] dx \\ & + \int_I^F [\ln Y - \ln Z/(F - I) - \ln \gamma_Z(x)] dx \\ & + \int_F^N [\ln Y - \ln((L - Z)/(N - F)) - \ln \gamma_H(x)] dx = 0 \quad (\text{B.11}) \end{aligned}$$

The equation then rearranged becomes:

$$\begin{aligned} \ln Y &= \int_{N-1}^I \left[\ln \left(\frac{A}{I - N + 1} \right) + \ln \gamma_A(x) \right] dx \\ & + \int_I^F \left[\ln \left(\frac{Z}{F - I} \right) + \ln \gamma_Z(x) \right] dx \\ & + \int_F^N \left[\ln \left(\frac{L - Z}{N - F} \right) + \ln \gamma_H(x) \right] dx \\ & = \int_{N-1}^I \ln \gamma_A(x) dx + \int_I^F \ln \gamma_Z(x) dx + \int_F^N \ln \gamma_H(x) dx \\ & + (I - N + 1) \ln \left(\frac{A}{I - N + 1} \right) + (F - I) \ln \left(\frac{Z}{F - I} \right) + (N - F) \ln \left(\frac{L - Z}{N - F} \right) \quad (\text{B.12}) \end{aligned}$$

The equation can then be simplified to show how aggregate output is a function of the different productivity levels of tasks combined with the supply of different factor inputs over the range of tasks those factor inputs cover.

$$Y = B \left(\frac{A}{I - N + 1} \right)^{I - N + 1} \left(\frac{Z}{F - I} \right)^{F - I} \left(\frac{L - Z}{N - F} \right)^{N - F} \quad (\text{B.13})$$

Then output Y in equilibrium is a product of the different relative productivities and the range of tasks that the different inputs of production cover.

$$B = \exp\left(\int_{N-1}^I \ln\gamma_{Ad}(x) + \int_I^F \ln\gamma_{Zd}(x) + \int_F^N \ln\gamma_{Hd}(x)\right) \quad (\text{B.14})$$

B is a collection of terms that encompass productivities of the different factor inputs.