

Artificial Intelligence: Where We Came From,
Where We Are Now, and Where We Are Going.

by

Guy-Warwick Evans
BSc, University of Victoria, 2013

A Master's Project Submitted in Partial Fulfillment of the
Requirements for the Degree of

MASTER OF SCIENCE

in the Department of Computer Science

© Guy-Warwick Evans, 2017
University of Victoria

All rights reserved. This project may not be reproduced in whole or in part, by
photocopying or other means, without permission of the author.

Artificial Intelligence: Where We Came From,
Where We Are Now, and Where We Are Going.

by

Guy-Warwick Evans
BSc, University of Victoria, 2013

Supervisory Committee

Dr. Bruce Kapron, Supervisor
(Department of Computer Science)

Dr. Sudhakar Ganti, Departmental Member
(Department of Computer Science)

Supervisory Committee

Dr. Bruce Kapron, Supervisor
(Department of Computer Science)

Dr. Sudhakar Ganti, Departmental Member
(Department of Computer Science)

ABSTRACT

From ancient myths to early advances in formal logic and mathematics, the story of artificial intelligence (AI) began centuries before the rise of modern computers. Today, modern AI has impacted nearly every area of human activity - from industries such as healthcare and transportation to science fiction media and popular culture. The story of artificial intelligence is far from over, with current trends and research suggesting large areas of impact in the future. This report examines three questions relating to AI: *where did AI come from*, *what is the current state of AI*, and *what does the future of AI look like*. A brief history of artificial intelligence is presented followed by a literature review and discussion of the impacts and trends of artificial intelligence research.

Contents

Supervisory Committee	ii
Abstract	iii
Contents	iv
List of Tables	vi
List of Figures	vii
1 Introduction	1
1.1 Report Structure	2
2 Where Did We Come From	4
2.0 Early Artificial Intelligence.....	4
2.1 The 19th and 20th Centuries	6
2.2 Modern Artificial Intelligence	9
2.3 Review: Where We Came From:.....	11
3 Where We Are Now	13
3.1 Impactful Artificial Intelligence Journals	13
3.2 Artificial Intelligence Research Funding.....	16
3.3 Review: Where We Are Now.....	18
4 Where Are We Going?	19
4.1 AI is Transforming Industries.....	19
4.2 Transportation	20
4.3 Medicine & Healthcare.....	22
4.3.1 Ethical Concerns Regarding AI in Healthcare.....	26
4.4 Military.....	27
4.4.1 Ethical Concerns Regarding AI in the Military	28
4.5 Displacement vs Replacement.....	29
4.6 General Artificial Intelligence.....	31

4.7 Review: Where Are We Going	33
5 Conclusion	34
Bibliography	37

List of Tables

Table 1 - Top 10 Artificial Intelligence journals ranked by impact factor	15
Table 2 - AI topics in AIME Conference 2017	25

List of Figures

Figure 1 - Yearly impact of top 3 AI journals	15
Figure 2 - Google Smart Car Sensors.....	21

ACKNOWLEDGMENTS:

I would like to thank:

Dr. Bruce Kapron for sparking my interest in computer science theory many years ago and teaching me an impressive number of classes (6) throughout my undergraduate and graduate schooling. I would also like to sincerely thank him for taking me on as an industrial master's student.

Dr. Sudhakar Ganti for being a fantastic teacher who always keeps his door open for his students, and always greets them with a warm welcome and guidance, and for stepping up to be my committee member last minute.

Lloyd Montgomery and Eirini Kalliamvakou for their friendship and guidance throughout my master's degree.

Wendy Beggs for her kindness, and her efforts in guiding graduate students through their time in graduate school.

Dr. Melanie Tory for teaching and guiding me in research during my research co-op and my first few terms as a master's student.

Dr. Daniela Damian for introducing me to the world of software development collaborative work and software engineering research, and for welcoming me and giving me a place in SEGAL for the first phase of my master's degree.

DEDICATION

I dedicate this report to my loving parents, Dianne & Warwick Evans.

Chapter 1

Introduction

In the modern world computers and the algorithms that govern them are seen everywhere; from the smartphones in our pockets, to the transportation systems we ride to work, to the computers that control our economy and banks. Many of these algorithms fall under the general umbrella of the field of artificial intelligence (AI). AI is a field that originally was founded by computer scientists in the 1950s but has since become a multidisciplinary field with applications in nearly every aspect of human life.

The field of AI was initially founded to answer the question: *is it possible to build a machine that has intelligence*, specifically a human level of intelligence. A necessary step in the pursuit of creating a machine intelligence was understanding the very nature of knowledge representation, reason, learning, perception, and problem solving. Through an understanding of these areas AI researchers discovered much narrower applications that a machine can perform and the field of artificial intelligence was expanded.

Today there are two broad classifications of artificial intelligence research, given to us in Peter Norvig's book *Artificial Intelligence: A modern Approach*. The first is *general artificial intelligence* or artificial general intelligence (AGI). AGI is concerned with the founding question of artificial intelligence - is it possible to create a machine that can perform the same intellectual tasks that a human being can. The second classification of AI research is known as *narrow artificial intelligence* or *weak AI*. Narrow

artificial intelligence is not concerned with creating a machine with human intelligence, and instead researchers in narrow AI solve specific narrow tasks.

Most of the artificial intelligence applications and research being done today falls into the narrow AI classification. Machine learning is one popular subfield of narrow AI that is concerned with finding methods to have computers learn how to solve a task or compute an answer without being explicitly programmed to do so. A popular type of machine learning seen today is known as supervised learning. In supervised learning a computer is given example inputs and their corresponding outputs, the computer learns a mapping rule based on these examples. When the computer is deployed in the wild and it is given a specific input, it then can compute a corresponding output.

This report aims to answer three questions regarding artificial intelligence:

1. *Where did we come from:* How did artificial intelligence begin?
2. *Where are we now:* What does artificial intelligence look like in the 21st century, what are the current trends in AI?
3. *Where are we going:* What does the future landscape of artificial intelligence look like? What kind of concerns are there?

This report will address those three questions providing a review of the history of artificial intelligence from antiquity to the 20th century and by reviewing the literature, government documents, and news releases of the 21st century.

1.1 Report Structure

Chapter 1 provides an introduction to modern artificial intelligence and the two categories of artificial intelligence research

Chapter 2 reviews the history of artificial intelligence from antiquity to the 20th century

Chapter 3 discusses the current trends in artificial intelligence research, the impact of top AI journals, and the corporate funding of AI start-ups

Chapter 4 provides a look at how artificial intelligence will impact three industries: transportation, healthcare, and the military. Chapter 4 also takes a look at the issue of unemployment due to automation, and finally chapter 4 looks at the issue of a future general artificial intelligence

Chapter 5 concludes the report

Chapter 2

Where Did We Come From

In a report discussing the impacts and trends of artificial intelligence it is pertinent to begin with asking the question *what is artificial intelligence?* This chapter aims to answer this question by beginning with a brief review of the history of artificial intelligence and ending with an introduction to modern artificial intelligence concepts. Through a look at the history of artificial intelligence this chapter will answer the question of *where we came from*.

2.0 Early Artificial Intelligence

Some researchers and historians trace the beginnings of the concept of artificial intelligence all the way back to ancient times where *thinking machines* and *artificial creatures* were used in myth and storytelling. The journalist and history writer Pamela McCorduck in her book “Machines Who Think” [40] references the greek myth of Talos, a man created out of bronze by the god Hephaestus to patrol and protect the beaches of Crete, along with the more famous story of Pandora and her box, also a creature created by Hephaestus, as two of the earliest examples of mythical thinking machines.

Some of the earliest known examples of non mythical machines that were built to exhibit some form of ‘tricked’ intelligence are described by Heron of Alexandria (10 AD-70 AD) in one of his famous works the *Pneumatica*. The *Pneumatica* describes several ideas for mechanical machines such as singing birds, and mechanical puppets [29]. While these machines were not created to

be thinking machines they were some of the earliest examples of mechanical machines built to perform specific tasks normally requiring intelligence.

Throughout the next several centuries the inventors and mechanics from many cultures continued to build mechanical toys and tools of increasing complexity, but never being able to capture real intelligence. In folklore and myth artificially intelligent beings continued to show up, such as the clay golems from Jewish myth. By the early 17th century the concept of complex machines and the topic of intelligence re-emerged with Descartes's writings on animals. Descartes explored the concept of the mind, body and soul. Descartes suggested that body works like a machine with material properties and that the mind or intelligence is non material, he used animals as an example of complex machines that have no reason or intelligence [40][21]. While Descartes's views might be outdated today, in fact animals not only have intelligence they may even have culture [34], his ideas of complex machines in the form of animals operating in our world was a novel one.

While the inventors of old attempted to build mechanical intelligence, the philosophers and mathematicians argued about the concept of logic and reasoning. Mathematical logic and formal reasoning were developed as a precursor to computers and modern artificial intelligence. By 1642 Blaise Pascal had invented what is credited as the first mechanical calculator, a 12 inch long mechanical brass box that could perform addition [12][47]. 30 years later in 1672 Gottfried Leibniz had improved on the previous models to create a mechanical calculator that could perform multiplication and division, now known as the Leibniz Calculator [47].

2.1 The 19th and 20th Centuries

Towards the end of the 19th century significant milestones in computational machines and logic were reached with notable mentions being George Boole who invented Boolean algebra (an important component of modern computers) along with Charles Babbage and Ada Lovelace who created what is credited as the first mechanical computer. The last interesting note from the 19th century that should be mentioned is the author Samuel Butler who was far before his time in suggesting that machines would one day gain consciousness and intelligence in his *The Book of The Machines* [16]. Samuel Butler also philosophized that machines might evolve in a similar manner as Darwinian evolution:

“..for as the vegetable kingdom was slowly developed from the mineral, and as, in like manner, the animal supervened upon the vegetable, so now, in these last few ages, an entirely new kingdom has sprung up of which we as yet have only seen what will one day be considered the antediluvian prototypes of the race.” - Samuel Butler [15]

Given that the 19th century was only the beginning of a mechanical revolution, Samuel Butler’s ideas were considered initially to be a joke. However, Mr. Butler was clear when he wrote the preface to his second edition that he was not joking.

“I regret that reviewers have in some cases been inclined to treat the chapters on Machines as an attempt to reduce Mr. Darwin's theory to an absurdity. Nothing could be further from my intention.” - Samuel Butler [16].

Samuel Butler’s ideas on conscious and evolving machines wouldn’t be revisited again until the mid 20th during the information technology revolution.

The beginning of the 20th century was an exciting time for mathematics and computer science and can be marked as a turning point for artificial intelligence progress. In 1913 Bertrand Russell and Alfred Whitehead published the *Principia Mathematica* which helped to create a foundation for mathematics and in particular formal logic. The Principia Mathematica in part attempted to describe a set of symbolic logic rules that could be used to derive all mathematical truths [30]. Twenty years later in 1931 Kurt Gödel published his incompleteness theorems which showed that the goals of the Principia Mathematica could actually never be reached by showing that for any sufficiently strong formal system either the system is inconsistent or there are mathematical truths that cannot be derived from that system. Gödel's incompleteness theorems have been used to argue against mechanistic theories of the mind, and as such occasionally are used to argue against the possibility of creating a truly intelligent artificial mind, however, Gödel himself acknowledged that his theories do not alone discredit the possibility of a machine mind [52].

In 1936 a landmark achievement was reached and the idea of the modern computer was born with Alan Turing publishing his paper *On Computable Numbers* [53]. Turing proposed a machine, called today a Turing Machine that is capable of computing anything that is mathematically computable. Turing machines use a tape and as such are capable of executing stored instructions (programs). Turing machines became the basis of what we know as the modern computer. Alan Turing contributed enormously to computer science in his short life and is now considered by some to be the father of computer science. Twenty six years after Alan Turing revolutionized computer science with his machines he wrote another paper titled *Computing Machinery and Intelligence* [54] where he proposed a test of a machine's ability to demonstrate intelligence now known as the *Turing Test*. The test is performed by having an evaluator interface through two text channels with both a computer programmed to answer questions as a human and an actual

human answering the questions. The goal of the test is for the evaluator to decide which text channel belonged to the human and which belong to the computer. Turing stated that if the evaluator was unable to distinguish between the channels then the computer is said to have passed the test. The Turing test today is a controversial test of intelligence, with most philosophers and scientists believing that it is not an adequate test for real intelligence. Some believe that the spirit of the Turing test is a good one and that if the Turing test is extended or modernized it remains a good test of artificial intelligence.

While the early computer scientists and mathematicians of the 20th century were building the foundations of what would later become modern computational theory the authors and novelists of the time were imagining what the world might look like in the future. In 1945, five years before Alan Turing proposed his Turing Test, the electrical engineer Vannevar Bush published his essay *As We May Think* [14] which was regarded as a visionary look at what the future of computers, artificial intelligence, and information science might be. *As We May Think* predicted much of the modern electronic landscape, including the use of personal computers as information tools assisting humans. Around the same time Isaac Asimov, considered one of the founders of modern science fiction, was busy writing and publishing many science fiction novels that had robotic characters and described futures where artificial intelligence played a key role. Notably Isaac Asimov's short story *Runaround* published in 1942 was the first time he made reference to his famous three laws of responsible robotics which he believed should govern artificially intelligent machines [6]:

1. *A robot may not injure a human being or, through inaction, allow a human being to come to harm.*
2. *A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.*

3. *A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.*

Isaac Asimov's three laws of robotics, along with his other works, have since been discussed at length in science fiction and academic circles. Some believe that the laws are sufficiently strong to govern artificially intelligent robots, others believe they are not. Susan Anderson, a professor emeritus at the University of Connecticut, who has contributed significantly to the area of machine ethics argues that Isaac Asimov did not intend for his laws to apply to truly artificially intelligent robots and that if they were to apply they would be immoral [3]. Nevertheless, in a time when computers were rare, slow, and large, Isaac Asimov's vision of a robotic future was novel.

2.2 Modern Artificial Intelligence

In 1956 a young computer scientist by the name of John McCarthy invited researchers from different computer science fields to a conference at Dartmouth College. In his invitation to these researchers he coined for the first time the term artificial intelligence and stated that the conference goal was:

"to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it." [33][39]

The goal of the organized conference was to bring together many different computer scientists with different expertise in order to discuss how to proceed with creating machine intelligence. John McCarthy would later on win the most prestigious prize in computer science research, a Turing award, for his

early contributions to artificial intelligence and to other computer science subfields. McCarthy is today considered to be the father of modern artificial intelligence.

The Dartmouth Conference would mark the beginning of a decade of large investment in artificial intelligence research with many promises for results and returns. Many researchers believed that machine intelligence was inevitable, and fast approaching. Isaac Asimov and other science fiction writers continued to write novels and short stories that pushed the boundaries of humanity's imagination for machine intelligence and the fantastic future that it would bring. Progress in other areas of computer science was swift as well, with computers becoming increasingly powerful and numerous.

By the late 1960s the initial promises of artificial intelligence researchers and theorists began to seem hollow. While after two decades progress in artificial intelligence research had been made in a number of subfields, the progress was slower than some expected. High initial enthusiasm lent itself to a large let down. In particular true machine intelligence seemed out of immediate grasp. The next period to begin in the history of artificial intelligence is referred to by some as the artificial intelligence winter (AI winter) where funding was reduced and enthusiasm dwindled.

As the AI winter continued progress in AI research was still made, in particular AI research began to focus less on creating true machine intelligence and instead on solving problems in other domains with AI tools. Health and medical sciences were one of the first research areas that saw interdisciplinary cooperation between computer scientists and medical researchers using artificial intelligence methods. The Dendral experiments conducted in the 1970s and 1980s brought together computer scientists, chemists, geneticists, and philosophers to build a computer system that could be used to interpret large quantities of high resolution mass spectral data [35][46]. In a special paper [46] based on a panel discussion held at the Artificial Intelligence in

Medicine Europe conference in 2007 it was discussed that the AI research community of the 1970s was fascinated by the new AI methods that were emerging as a result of AI researchers working on applications in the medical sciences domain. The panel also noted that in 1978 the leading journal on artificial intelligence at the time devoted a special issue specifically for medical AI research. Furthermore, by 1980 when the American Association for Artificial Intelligence was formed a special subgroup on medical applications was also created.

During the AI winter the first major technology laboratories supporting AI were founded at the Massachusetts Institute of Technology and Carnegie Tech, along with laboratories at Stanford and Edinburgh [13]. Many different AI subfields were created and the concept of modern artificial intelligence was fleshed out. There doesn't seem to be any consensus on when the AI winter began and when it ended, some even argue that with all the progress that had been made there was no winter. One can argue that certainly by 1997 when IBM's Deep Blue [23] computer beat a grandmaster chess player for the first time that if there was an AI winter, it had been long over.

2.3 Review: Where We Came From

This chapter reviewed some of the highlights of the history of artificial intelligence, and from those highlights we can now gain a picture of *where we came from*.

From as far back as antiquity humans imagined the use of artificially intelligent machines as protectors, helpers, entertainers, and existential threats. While our imaginations were large the early progress in artificial intelligence was slow, with the foundations of mathematics and computing needing to be built first. For centuries as our knowledge of mathematics slowly lead us towards the first computers we continued to explore the idea of

artificially intelligent beings in philosophy and myth. Given our large imaginations it is no surprise that once the first computers were being built the enthusiasm for artificial intelligence exploded both within the scientific community and the science-fiction community.

However, when computers facilitated artificial intelligence research and experimentation, academics began to realize that a successful general AI might be a long way away off, or even unreachable. As enthusiasm for general AI dwindled among researchers the enthusiasm for narrow AI began to grow. Academics and investors were realizing that the impact that narrow AI tools could have to interdisciplinary fields was large. Artificial intelligence expanded and grew into a large scientific endeavour with roots in many different research areas.

Finally a full picture of *where we came from* has formed. While humanity's initial dream for artificial intelligence was to build a true machine intelligence, the field of artificial intelligence has grown into a much larger and much more applied field. The goal of creating that true machine intelligence still exists for some, but it is no longer the only goal of AI with other avenues of inquiry into narrow AI being opened up.

Chapter 3

Where We Are Now

In the previous chapter I outlined a brief history of artificial intelligence. From this history an understanding of *where we came from* was reached. The previous chapter took us to the end of the 20th century when artificial intelligence had grown from a discipline with one question at its core: *can we create machine intelligence* to a multi-discipline field with many narrow questions. The next question this report will address is *where are we now?* To understand what the current artificial intelligence landscape looks like this chapter will look at some of the trends and research in AI from the beginning of the 21st century to today.

3.1 Impactful Artificial Intelligence Journals

I began by investigating the highest impact conferences and journals in AI today. The Journal Citation Reports (JCR) [32] is an annual publication by Thomson Reuters that can be accessed through the Web of Sciences [57] portal. The JCR provides information about journals and conferences along with the impact factors of those journals. The impact factor of an academic journal is one measure of looking at the academic importance of that journal. Impact factor of a journal is the number of citations received in that year by articles published in that journal during the two preceding years, divided by the total number of articles published in that journal during the two preceding years. Using the JCR I looked at the top ten journals in 2015 (the most recent data available) that were categorized as ‘Computer Science: Artificial Intelligence’ and ranked by impact factor. These journals are listed in table 1.

Immediately looking at the journals in table 1 it is evident just how varied artificial intelligence research is, with so many top journals devoted to applied artificial intelligence. However, to understand the impact of these journals it is useful to compare the impact factors against the impact factors of other journals in computer science. The *IEEE Communications Surveys and Tutorials* journal has an impact factor of 9.22 according to the JCR search, it is the highest impact journal in computer science. We can immediately see that the impact of the top AI journals is within a comparable range to the top computer science journal. In fact, the top AI journal, IEEE Transactions on Fuzzy systems, is ranked as the number two computer science journal according to the JCR impact rating. From these results we can see that artificial intelligence research has a large presence among contemporary computer science research.

The JCR impact rating and database also allows us to look back in time to see how the impact of AI journals have changed since the beginning of the 21st century. Figure 1 shows the change in JCR impact for the top three (in 2015) AI journals.

From figure 1 we can see how the impact rating of the three top AI journals has improved since 1998. From this we can conclude that artificial intelligence research impact appears to be on the rise and its presence in the academic community is increasing every year.

	Journal	Citations	Impact Factor
1	IEEE TRANSACTIONS ON FUZZY SYSTEMS	9,220	6.701
2	International Journal of Neural Systems	1,235	6.085
3	IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE	31,757	6.077
4	IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION	7,999	5.908
5	INTEGRATED COMPUTER-AIDED ENGINEERING	533	4.981
6	IEEE Transactions on Cybernetics	2,246	4.943
7	IEEE Transactions on Neural Networks and Learning Systems	12,919	4.854
8	MEDICAL IMAGE ANALYSIS	4,764	4.565
9	Information Fusion	1,637	4.353
10	INTERNATIONAL JOURNAL OF COMPUTER VISION	11,407	4.27

Table 1: Top 10 Artificial Intelligence journals ranked by impact factor

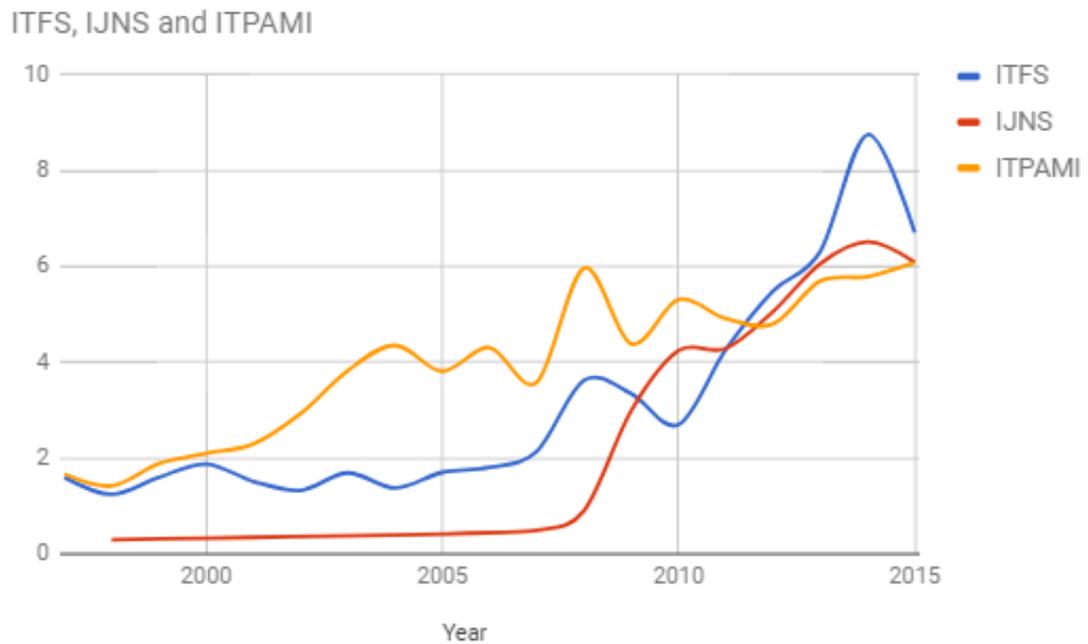


Figure 1: Yearly impact of top 3 AI journals [32]

3.2 Artificial Intelligence Research Funding

The academic world is not the only place that artificial intelligence is currently seeing increasing impact. The corporate and business world has also seen massive AI successes in the first quarter of the 21st century with companies such as Google and Facebook utilizing sophisticated searching and machine learning algorithms in multiple products. To get a sense of the current enthusiasm and impact of AI on the corporate world I investigated the literature on AI funding.

CB Insights is a business that aggregates and analyzes market and internal data for their customers. In 2017 CB insights published a recap of the funding received in artificial intelligence related deals in the five years leading up to 2016 [17]. CB insights found that deals had reached a five year high with 160 deals in 2012 and 658 in 2016. CB insights also found that the amount of dollars invested in AI businesses rose by 60% over the five year period. CB insights found that the final quarter of 2016 was the most active quarter in investments in the five year time with huge investments in a variety of applied AI start-ups.

According to the Pardee Center for International Futures background report on Artificial intelligence [50] written for the U.S. National Intelligence council under the Obama administration in 2016 the global artificial intelligence market in 2016 is valued at \$126 Billion US. The Analysis Group, a company that provides economic and financial consulting, published their own report in 2016 that also points out that much of the AI industry is currently in the form of start-ups. From 2014 to 2015 the eight leading tech firms (Google, Microsoft, Apple, Amazon, IBM, Yahoo, Facebook, and Twitter) made 26 acquisitions of artificial intelligence related start-ups valued at an estimated \$5 billion US [50].

In addition to private corporate funding, governments are also increasing their AI research funding. On March 22nd 2017 the Government of Canada announced funding for a Pan-Canadian artificial intelligence strategy for research and talent to cement Canada's position as a world leader in AI. 125 million funding will be allocated to attract the top AI academic talent to Canada [45]. The Government of the United Kingdom also announced a £20 million investment into artificial intelligence research in 2016 [49].

From these results it is evident artificial intelligence research has seen increasingly large investments in the first quarter of the 21st century with governments and high tech firms considering narrow AI research to be an important investment. While narrow AI is seeing huge investment, humanity has not given up on the original goal of creating a general artificial intelligence. Several research institutions are currently working toward that goal such as the Machine Intelligence Research Institute [38], Numenta [43], and the Helen Wills Neuroscience Institute [28]. Open source collaborative projects for general artificial intelligence also exist, such as the OpenCog [44] project which is directly aimed at creating a beneficial general artificial intelligence.

Artificial intelligence has continued to have an impact on popular culture with AI appearing in everything from music videos, novels, movies, and video games. AI characters in popular culture are capturing the imagination of many such as TARS from *Interstellar*, and GLaDOS in the video game *Portal*. Wikipedia alone lists 49 films created in the first quarter of the 21st century featuring an artificial intelligence [36].

3.3 Review: Where We Are Now

This chapter reviewed the progress that has been made, and the influence of, artificial intelligence in the first quarter of the 21st century. I now return to the question of *where we are now*. Continuing the trends from the previous century enormous progress and investments are being made in narrow artificial intelligence research. The benefit of using artificial intelligence tools and methods to solve specific problems in different industries and research areas has been noticed and work is being done to continue this progress. Despite continued work general artificial intelligence still remains out of grasp but continues to capture the imaginations of science fiction content creators and influence popular culture.

Chapter 4

Where Are We Going?

The previous two chapters answered the questions of *where we came from* and *where are we now* in regards to the field of artificial intelligence. In this chapter I will address the question of *where are we going* and present some of the expected trends and consequences of artificial intelligence research in the immediate and distant future. This chapter will also discuss some of the recent calls for caution around general artificial intelligence made by leading scientists.

4.1 AI is Transforming Industries

On January 25th, 2017 Andrew Ng an adjunct professor at the University of Stanford, former chief scientist at Baidu, and founder of the Google Brain deep learning research project gave a talk titled *Artificial Intelligence is the New Electricity* [42]. In his talk Andrew argued that artificial intelligence is poised to transform many different industries. Andrew argued that the IT industry has already been transformed by AI as seen by the huge successes of tech companies such as Google, Amazon, and Microsoft. Andrew mentions the transportation, logistics, and healthcare industries all as examples of industries that are currently being transformed or poised to be transformed in the coming immediate future. In fact, Andrew states:

“Just as electricity transformed almost everything 100 years ago, today I have a hard time thinking of an industry that I don’t think AI will transform” - Andrew Ng [42]

A number of recent studies have been conducted that support Andrew Ng’s hypothesis that artificial intelligence is poised to transform the world. Most notably, the Pardee Report prepared for the National Intelligence Council in 2016 by Andrew Scott and Barry Hughes [50] is a lengthy review of what the future of artificial intelligence might look like in the United States and globally. The report makes the clear point that artificial intelligence is poised for rapid advancement and that AI has the potential to transform many areas of human activity. Additionally, the Pardee report predicts that labour will be impacted by the automation of jobs, particularly in the manufacturing and service sectors.

4.2 Transportation

When it comes artificial intelligence transforming industries the first industry that comes to most people's mind is the transportation industry. There has been lots of recent enthusiasm among tech blogs and the media around the success of self-driving technology. On April 4th 2016 a convoy of six semi-automated trucks demonstrated that success by driving from multiple factories in different countries to Rotterdam, Netherlands using self-driving technology. The trucks had human drivers behind the wheel for emergencies but were otherwise completely self-driving [25]. Later on that year Uber’s self driving truck, built by the start-up company Otto, successfully made a 120 mile drive from Fort Collins to Colorado Springs with a delivery of 50,000 cans of beer [22].

Google's smart driving technology has also seen enormous recent success with the Google smart car fleet have successfully driven more than 3 million miles. While it took the Google smart car fleet six years to drive the first 1 million miles, it only took two more years to drive the next 3 million - showing a great improvement in the technology [24]. Figure 2 created by Google shows how well the Google smart car sensors are, being able to sense and predict not only the road and the immediate pedestrians but also a cyclists left-turn hand signal and distant cars.

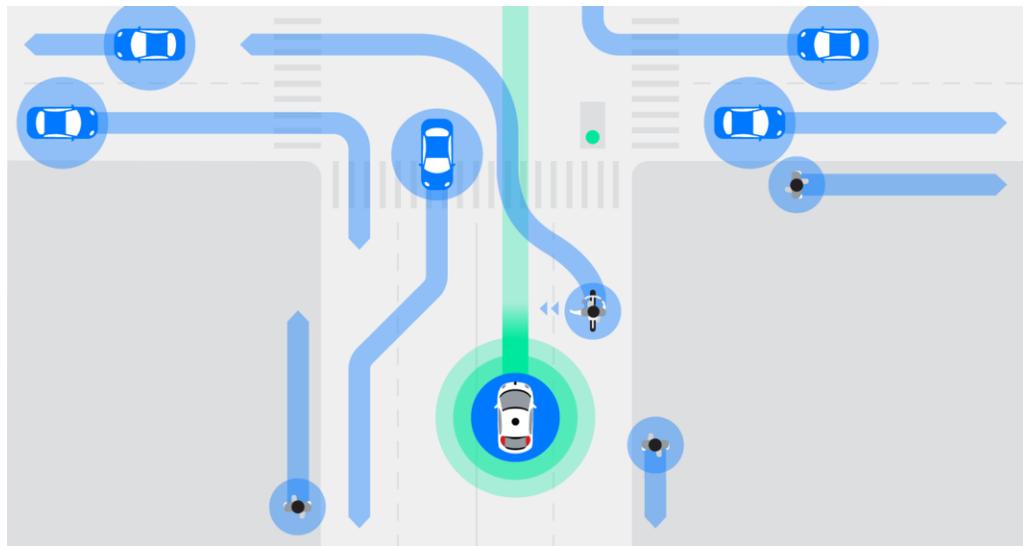


Figure 2 - Google Smart Car Sensors

With 94% of car crashes in the United States involving some form of human error according to the 2016 National Motor Vehicle Crash Causation survey [41] the motivation to replace conventional driving with smart driving technology is there. The Pardee report estimates that autonomous vehicles could reduce traffic accidents by up to 90%, saving the US economy \$400 billion annually. However, adopting self driving technology would be a very large culture and societal change with many people hesitant to trust an artificial intelligence driving for them.

One concern for the adoption of smart driving technology is the potential for mass unemployment of transportation occupations. The American Trucking association states that there are over 3.5 million truck drivers employed in the united states [2] and the US Bureau of Labor Statistics states that there were 233,70 employed taxi drivers in the US in 2014 [55]. Both truckers and taxi cab drivers, along with many other people employed in transportation services, may stand to lose their jobs if self driving technology goes mainstream.

Despite the concerns about the impact that self driving technology might have on transportation labour, it is clear that the technology is not slowing down with companies such as Google and Uber investing large amounts of money. Self driving technology is not the only area of transportation that artificial intelligence is likely to impact. Other areas include road maintenance, road safety and planning, traffic prediction, control systems, and construction. Artificial intelligence will shape transportation in the immediate future.

4.3 Medicine & Healthcare

Another industry frequently mentioned as being on the cusp of large changes due to artificial intelligence is healthcare and medicine. As mentioned in chapter 2 artificial intelligence use in medicine has a long history, dating back to the 1970s and 1980s when enthusiasm for AI use in medicine was high. In the research community artificial intelligence use in medicine is still alive and well today with several high impact journals and conferences, namely: SIIM Conference on Machine Intelligence in Medical Imaging, the MICCAI Medical Imaging Society journal, and the Artificial Intelligence in Medicine journal.

In 2007 at the Artificial Intelligence in Medicine conference in Amsterdam a panel discussion was held titled “The coming age of artificial intelligence in medicine” [46]. In the panel six speakers including the chair

discussed the previous trends and impact of artificial intelligence in medicine and discussed what the future of AI in medicine might look like. One of the speakers, Dr. Edward Shortliffe, a physician and computer scientist and regarded as one of the founders of biomedical informatics assessed that artificial intelligence in medicine has progressed quietly in the 21st century and under various narrow applications not always labeled as artificial intelligence. Dr. Shortliffe outlined several key problems that he saw as needing to be tackled to progress the future of AI and medicine. One such problem is the need to educate more professionals in the interdisciplinary nature of biomedical informatics. The panel members went on to discuss various aspects of medical informatics that need improvement, such as confidentiality in medical applications and informatics systems that model physician workflow in order to work in conjunction with the physicians. The closing remarks of the panel discussed the rising amount of data in clinical applications and the opportunity that artificial intelligence methods has to use that data to benefit health care moving forward.

One specific area of healthcare that may see great changes due to artificial intelligence is the discipline of radiology. Katie Chockley and Ezekiel Emanuel in their paper *The End of Radiology? Three Threats to the Future Practice of Radiology* outline what they think are the three greatest threats to radiology, the ultimate threat among them being the rise of artificial intelligence [19]. Radiology is a discipline that operates on the interpretation of medical images to lead to a diagnosis. With the increasingly good imaging capabilities of modern medical systems the amount of medical image data is increasingly enormously, which increased the workload for a radiologist who needs to use the images for a diagnosis. Chockley et al. point out that humans are susceptible to fatigue and other distractions such as emotions, whereas machines are not. Additionally machines are becoming better and better at interpreting images - and in some cases surpassing the ability of a trained radiologist [10]. Jha et al. in an editorial for the magazine *Innovations in*

Healthcare Delivery points out that in order for radiologists to avoid being replaced entirely they should begin training as information specialists [31].

It is clear from the literature that there is no consensus on just how much, and how quickly artificial intelligence will revolutionize medicine. Radiology might stand poised to change more rapidly than other areas of healthcare, however, artificial intelligence is being applied in a very large number of areas, and table 2 shows just some of the categories that are being accepted in the Artificial Intelligence in Medicine conference in 2017. Seasoned medical informaticians such as the ones from the “The coming age of artificial intelligence in medicine” panel might be hesitant to be too enthusiastic about artificial intelligence as they were witness to the hype and letdowns of the 1980s. However, as AI knowledge and methods improve it is clear there are a large number of areas both in clinical practice and in biomedical research that will be impacted.

Big Data Analytics	Biomedical Knowledge Acquisition and Knowledge Management
Clinical Decision Support Systems	AI methods in Telemedicine and eHealth
Behavior Medicine	Patient Engagement Support (Personal Health Record)
Machine Learning, Knowledge Discovery and Data Mining	Case-based Reasoning in Biomedicine
Biomedical Ontologies and Terminologies	Document Classification and Information Retrieval
Bayesian Networks and Reasoning Under Uncertainty	Biomedical Imaging and Signal Processing
Temporal and Spatial Representation and Reasoning	Visual Analytics in Biomedicine
Computerized Clinical Practice Guidelines and Protocols	Natural Language Processing
Fuzzy Logic	Healthcare Process and Workflow Management
AI solutions for Ambient Assisted Living	

Table 2- AI topics in AIME Conference 2017 [20]

4.3.1 Ethical Concerns Regarding AI in Healthcare

Most of the ethical concerns regarding AI use in healthcare are centered on the issue of artificial intelligence use in *diagnosis* and *treatment*. For example, an AI system that might replace a radiologist and diagnose a patient based on medical imaging. Another example could be a system that suggests a particular treatment course for a patient, or even a system that performs surgery on a patient. The central issue of these AI systems is in the case of a wrong diagnosis, treatment, or an incorrect surgical procedure who is held responsible? The hospital, or the company who built the system, or the engineers who designed the system? These answers are not immediately clear and likely will not be clear until challenged in a particular legal system.

Another ethical concern that is often raised regarding AI in medicine is the concern of confidentiality. AI systems by their very nature require large amounts of data in order to perform well. Patient data is some of the most sensitive and personal data in the information world and the way an AI system uses a patient's data is a large concern, particularly because most AI systems would be required to not only use the patient's data that they are treating, but to have access to a database of other patient data in order to decision about a correct treatment procedure.

Another ethical concern is a loss of humanity in the healthcare if machines replace certain kinds of doctors. Is it important to have a human doctor, who is susceptible to fatigue, emotions and bias, to provide care? Or can a machine provide the same care without the drawbacks.

These ethical concerns don't have clear answers and as AI delivered healthcare is in its infancy they have not been explored in depth.

4.4 Military

The final industry that I will review in this report is the military. In particular the United States military's funding and research into artificial intelligence warfare.

In 2015 US Deputy Defense Secretary Robert Work said they planned on having at least \$12 billion funding set aside by 2017 for artificial intelligence weapon technology [51]. This technology will include autonomous weapons and deep learning machines that focus on human-machine collaboration in combat. The Defense Advanced Research Project Agency (DARPA) has long funded artificial intelligence and robotics research and runs a yearly robotics challenge known as the DARPA Robotics Challenge to both recruit potentially skilled talent and to advance the robotics industry.

What does artificial intelligence use in the military entail? Autonomous weapons are the first thing to come to mind. Autonomous weapons are artificially intelligence weapons that make kill decisions based on their programming. Systems that operate in logistics and surveillance are another application for AI in the military. Artificial intelligence can also be deployed to protect against cyber attacks, or conduct cyber attacks against rival powers.

Ayoub et al. argue that artificial intelligence will in the near future have a profound impact on the conduct of military strategy [8]. Specifically they argue that narrow artificial intelligence in the immediate future will utilize large sums of data in military strategic decision making and make decisions based on that data faster than a human is capable of making. Humans are susceptible to fatigue and emotional decision making and as such Ayoub et al argue that an artificial intelligent system that is sufficiently trained is better positioned to make strategic military decisions in certain contexts. The first nations to use an AI successfully to make military decisions could have an upperhand in any kind of military conflict and as such this has the potential to shift the balance of power between nations.

In 2015 over 1000 leading artificial intelligence and technology experts including famed physicist Stephen Hawking, Apple Co-founder Steve Wozniak, and Tesla CEO Elon Musk, signed an open letter that urges a ban on warfare artificial intelligence and autonomous weapons [7]. The letter warns of a military artificial intelligence arms race beginning and that the practical deployment of autonomous weapons is feasible within years and that their deployment will constitute a third revolution in warfare, after gunpowder, and nuclear arms.

Edward Geist argues that an artificial intelligence arms race has already begun and it is too late to prevent it from continuing [27] and that instead a culture of security should be cultivated to manage the AI arms race. Dr. Geist also argues that autonomous weapons that can kill are a foregone conclusion and that it will not be possible to come to an international agreement to ban them.

While military institutions are reasonably secret about the full extent of their research into new weaponry and tactics it is clear that artificial intelligence is being used today, and will continue to be used an even greater amount in the future by any modern military power.

4.4.1 Ethical Concerns Regarding AI in the Military

There are a number of ethical concerns regarding the use of artificial intelligence in the military, specifically the use of lethal autonomous weapon systems (LAWS). LAWS are machines that could identify and attack a target without any kind of human intervention. LAWS have been a major discussion point among ethicists and at the United Nations. The issue of LAWS was first brought to the international community's attention by a report published by the Human Rights Watch titled: *Losing Humanity: The Case Against Killer Robots* [37]. The report predicts that 'killer robots' could select and engage targets without human intervention within 20 to 30 years. The report argues

that killer robots would not be consistent with international humanitarian law and pre-emptive prohibition on their development and use is needed.

The issue of LAWS being ethical to build and use is a matter of ethical frame point. A utilitarian might see LAWS as ethical and may argue that machines make less mistakes and as such the utility of using them outweighs the utility lost. Peter Asaro argued at the 2014 meeting of experts on LAWS hosted by the UN that if we dehumanize warfare by deploying LAWS then there will be a loss of justice, accountability, and responsibility. Dr. Asaro also argues that allowing an automatic process to decide to take human life, diminishes the value of human life [5].

The UN is taking LAWS seriously, having hosted several panels of military, legal, and ethical experts to discuss their use and to discuss if a ban should be put into place, and if a ban would be effective to prevent their use. From reviewing the published UN papers on the use of LAWS [9] it seems clear that the consensus among ethicists is that the use of LAWS is not ethically consistent with the UN values, however, it also seems clear that the use of LAWS are inevitable.

4.5 Displacement vs Replacement

An issue that was brought up in the previous sections is the potential for artificial intelligence systems to replace humans in jobs. Truckers were mentioned as being particularly vulnerable if self driving technology goes mainstream, and radiologists were mentioned as another vulnerable profession. There are many other jobs that stand to be replaced by artificially intelligent machines such as: fast food workers, retail salespersons, security guards, and receptionists. Even delivery drivers could be replaced by autonomous drones [1].

In 2016, the head of the Canadian government economic growth advisory council, Dominic Barton, said that about 40% of existing Canadian jobs could disappear over the next decade due to automation [18]. Dr. Judy Wajcman, a professor of sociology at the Long School of Economics and Political Science, conducted a review of several leading books on the effect of automation and artificial intelligence on jobs [56]. In her review Dr. Wajcman found that most of the literature predicted that AI would cause unprecedented job loss. A study conducted in 2015 by Frey et al. estimates that 47% of total US employment is in the high risk category when it comes to replacement by computerization [26], in fact this is the study that Dr. Wajcman points out is used predominantly in the literature to support the claim that AI will cause unprecedented job loss, and Dr. Wajcman further points out in her review this study's methodology has been brought into question by another study by Arntz et al [4]. Arntz et al. argue that the Frey et al. study overestimated job automatability. Arntz et al. conducted their own review of automatability and found that across the 21 countries they studied only 9% of jobs are automatable.

To the best of my knowledge there is no real agreement among the literature just how much artificial intelligence will impact jobs. Historically technology has even ended up creating more jobs than it has destroyed. For example, during the Industrial Revolution when weaving was being replaced by machines it was thought at first that weavers would become unemployed, but instead they were employed to tend the machines. With the increase in weaving productivity the cost of cloth plummeted which increased the high elastic demand, causing even more people to be employed working the machines than were originally employed weaving [11]. James Bessen argues that it is only in the case of complete automation that there is a net loss of jobs, and historically in the case of partial automation of a profession then there is often a net increase of jobs [11]. Bessen rejects computer automation as a

significant source of *overall* job loss, but notes that some occupations that are automated shift employment to other jobs that require different skills.

The shift from one job to another job brings up the question of *displacement vs. replacement*. Will artificial intelligence cause humans to be replaced in their jobs and therefore cause mass unemployment? Or will artificial intelligence cause humans to be displaced to other jobs. Returning to the case of a radiologist being replaced by a machine, if the radiologist instead trains as an information specialist then that individual may still have a job, and instead of being replaced by a machine they will have been displaced. Moving forward into the future it will be necessary for researchers and governments to be aware of the effects of automation and artificial intelligence on various disciplines, so that governments can take early action to facilitate job displacement in order to prevent job replacement, and therefore mass unemployment and inequality.

4.6 General Artificial Intelligence

So far in chapter 4 I have discussed narrow artificial intelligence, which leaves the question of *where are we going* in regards to general artificial intelligence. In chapter 3 I briefly reviewed the ongoing research into creating a general artificial intelligence and concluded that while it is still an active research field, we are still many years away from successfully creating a general AI. In chapter 2 I reviewed some of the earliest human explorations of general artificial intelligence and showed that the question of general artificial intelligence was the founding question of the field of artificial intelligence.

So what does general artificial intelligence look like in the future? The answer to this question is simply, no one knows for sure. Futuristic landscapes with general AI have been explored since the early 19th century in the form of

science fiction novels, and continue to be explored in popular culture today in movies such as Terminator and science fiction novels such as Cloud Atlas.

Some futurists and researchers have attempted to answer the question of future general AI. Ray Kurzweil a popular futurist, computer scientist, and author has written several novels where he estimates that machines will be intelligent by 2029 and that not long after due to the incredible exponential pace of technological innovation a superintelligence emerge in an event known as the technological singularity (or simply, the singularity) [48]. The singularity is an event triggered by runaway technological growth that will result in massive changes to human life and civilization. The premise of the singularity is that a general artificial intelligence will know how to improve itself, and will continue to improve itself in a loop fashion until the AI has transformed into a super intelligent AI.

The concept of a singularity and a superintelligence are controversial in the academic community, particularly the timescale that Ray Kurzweil proposes. However the idea that a general artificial intelligence might one day be created is one that many academics do take seriously. Famed physicist Stephen Hawking and tesla CEO Elon Musk have made warnings in the past that humanity must proceed with caution when creating a general artificial intelligence due to the risk that a general AI might not be beneficial to humanity. AI researchers Stuart Russel and Peter Norvig in their textbook *Artificial Intelligence: A Modern Approach* discuss the possibility that due to an artificial intelligence's learning function an AI system might evolve into a system with unintended behaviour. This unintended behaviour of an AI is a concept that has been explored at length in science fiction with the most famous example being the AI Skynet that was developed in the movies Terminator to improve mankind, instead Skynet evolved into a machine that wanted to eradicate mankind.

While researchers work to solve that ancient artificial intelligence question and create a true machine intelligence there is a need for other

researchers, and science fiction writers, and policy makers to continue working on how to make sure that if humanity ever has the ability to create a general AI, we can do so in a responsible way that does not pose an existential threat to ourselves.

4.7 Review: Where Are We Going

In this chapter I addressed the question of *where are we going* with artificial research and what kind of impacts and trends will that research have. I reviewed the effect that narrow artificial intelligence will have on three industries: transportation, healthcare, and the military. I also presented some of the basic ethical concerns that deployment of AI methods might raise. I discussed the need for governments to facilitate job displacement in order to prevent job replacement in the event of automation of jobs. Finally, I reviewed some of the concerns about general artificial intelligence in the future and the existential threat that it might pose, along with the need to consider how to proceed with general artificial research in a responsible manner.

Chapter 5

Conclusion

In this report I aimed to answer three questions about the past, present, and future of artificial intelligence:

1. *Where did we come from*
2. *Where are we now*
3. *Where are we going*

In Chapter 2 I reviewed the history of early artificial intelligence from antiquity leading up to the 20th century when the field of modern artificial intelligence was founded. From this history we saw that the humanity had dreams of intelligent machines all the way back in ancient times, and we saw how early advancements in mathematics and logic lead to the creation of the modern computer. We saw how the modern computer gave new hope to that early dream of machine intelligence and sparked massive enthusiasm in artificial intelligence research. In trying to understand the concept of intelligence and trying to create that intelligence the field split into two, one in pursuit of general AI and the other in pursuit of Narrow AI. From that pursuit of narrow AI computer scientists and other interdisciplinary scientists found enormous use and application of AI methods.

In Chapter 3 I reviewed what the current state of artificial intelligence research looks like. We saw that throughout the first quarter of the 21st

century artificial intelligence research has been gaining impact in the research community. We saw that funding in artificial intelligence research has been increasingly steadily, particularly in the last five years. Finally, we saw that artificial intelligence had a huge effect on popular culture.

In Chapter 4 I reviewed three industries that are poised to be impacted by artificial intelligence in the near future: transportation, healthcare, and the military. I presented the current consensus from the literature in these industries and discussed some of the ethical concerns surrounding increased AI utilization in these industries. I discussed the potential for mass unemployment due to AI and the need for governments and industries to consider how to avoid job replacement, by facilitating job displacement.

This work has lead me to the answer my research questions as follows:

1. *Where did we come from:* The field of artificial intelligence was born out of a single question: can we create a true machine intelligence? While humanity's initial dream for artificial intelligence was to build that machine intelligence, the field of artificial intelligence has grown into a much larger and much more applied field with many areas of pursuits and many applications.
2. *Where are we now:* Artificial intelligence research is varied and becoming increasingly more impactful with large amounts of money from both private and public institutions being invested. Artificial intelligence continues to capture the imaginations of humanity as demonstrated by its use in popular culture.
3. *Where are we going:* Artificial intelligence is a fact of the future. Its effects will be seen in potential every industry and those effects might be profound. The use of intelligent machines could be enormously beneficial to humanity, from improving our healthcare system to preventing vehicle deaths. However, intelligent machines could also

pose significant risk, from autonomous machines making kill decisions, to massive unemployment, to the potential of a general AI existential threat.

Bibliography

1. Amazon Prime Air. Amazon. Accessed 05-10-7.
<https://www.amazon.com/Amazon-Prime-Air/b?node=8037720011>
2. American Trucking Associations. "Reports, Trends & Statistics".
Accessed 5-10-07.
http://www.trucking.org/News_and_Information_Reports_Industry_Data.aspx
3. Anderson, Susan Leigh. "Asimov's "three laws of robotics" and machine metaethics." *Ai & Society* 22.4 (2008): 477-493.
4. Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. "The risk of automation for jobs in OECD countries: A comparative analysis." *OECD Social, Employment, and Migration Working Papers* 189 (2016): 0_1.
5. Asaro, Peter. School of Media. Ethics of LAWS Talk. https://unoda-web.s3-accelerate.amazonaws.com/wp-content/uploads/assets/media/79F6199F74DC824CC1257CD8005DC92F/file/Asaro_LAWS_ethical_2014.pdf
6. Asimov, Isaac. *I, robot*. Vol. 1. Spectra, 2004.

7. "Autonomous Weapons: An Open Letter From AI & Robotics Researchers." Future of Life Institute. 06-28-2016.
<https://futureoflife.org/open-letter-autonomous-weapons/>
8. Ayoub, Kareem, and Kenneth Payne. "Strategy in the Age of Artificial Intelligence." *Journal of Strategic Studies* 39.5-6 (2016): 793-819.
9. Background on Lethal Autonomous Weapons Systems. United Nations Office for Disarmament Affairs. Accessed 05-10-07.
<https://www.un.org/disarmament/geneva/ccw/background-on-lethal-autonomous-weapons-systems/>
10. Beck, Andrew H., et al. "Systematic analysis of breast cancer morphology uncovers stromal features associated with survival." *Science translational medicine* 3.108 (2011): 108ra113-108ra113.
11. Bessen, James E. "How computer automation affects occupations: Technology, jobs, and skills." (2016).
12. Blaise Pascal and the First Calculator By: Karwatka, Dennis. *Tech Directions*. Nov2004, Vol. 64 Issue 4, p10-10. 1p. 2 Black and White Photographs.
13. Buchanan, Bruce G. "A (very) brief history of artificial intelligence." *Ai Magazine* 26.4 (2005): 53.
14. Bush, Vannevar. "As we may think." *The atlantic monthly* 176.1 (1945): 101-108.

15. Butler, Samuel. "Darwin among the machines." (1973). Accessible at: <http://www.gutenberg.org/files/6173/6173-h/6173-h.htm>
16. Butler, Samuel. "Erewhon." (1973) Chapter: The Book of the Machines. Accessible at: <http://www.gutenberg.org/files/1906/1906-h/1906-h.htm>
17. "The 2016 AI Recap: Startups See Record High In Deals And Funding". CB INSIGHTS. January 19, 2017. <https://www.cbinsights.com/>
18. "Canada must rethink skills training as automation eliminates jobs: Barton." The Canadian Press. 02-06-07. <http://ipolitics.ca/2017/02/06/canada-must-rethink-skills-training-as-automation-eliminates-jobs-barton/>
19. Chockley, Katie, and Ezekiel Emanuel. "The End of Radiology? Three Threats to the Future Practice of Radiology." Journal of the American College of Radiology 13.12 (2016): 1415-1420.
20. Conference on Artificial Intelligence in Medicine. Topics for papers. Access 05-10-07. <http://aime17.aimedicine.info/call-for-papers-submission.html>
21. Cottingham, John. "A Brute to the Brutes?: Descartes' Treatment of Animals." Philosophy 53.206 (1978): 551-559.
22. Davies, Alex. Uber's Self-Driving Truck Makes Its First Delivery: 50,000 Beers. WIRED. 10-25-16. <https://www.wired.com/2016/10/ubers-self-driving-truck-makes-first-delivery-50000-beers/>

23. "Deep Blue". International Business Machines (IBM). 05-10-07.
<http://www-03.ibm.com/ibm/history/ibm100/us/en/icons/deepblue/>
24. WAYMO. Google Smart Car. Accessed 05-10-07.
<https://www.google.com/selfdrivingcar/>
25. France-Presse, Agence. "Convoy of self-driving trucks completes first European cross-border trip."
<https://www.theguardian.com/technology/2016/apr/07/convoy-self-driving-trucks-completes-first-european-cross-border-trip>
26. Frey, Carl Benedikt, and Michael A. Osborne. "The future of employment: how susceptible are jobs to computerisation?." *Technological Forecasting and Social Change* 114 (2017): 254-280.
27. Geist, Edward Moore. "It's already too late to stop the AI arms race— We must manage it instead." *Bulletin of the Atomic Scientists* 72.5 (2016): 318-321.
28. Helen Wills Neuroscience Institute. University of California. <http://neuroscience.berkeley.edu/>
29. "Heron of Alexandria." *Encyclopedia Britannica*. *Encyclopedia Britannica*, 1-24-017. Accessed: 6-10-2017.
<https://www.britannica.com/biography/Heron-of-Alexandria>
30. Irvine, Andrew David, "Principia Mathematica", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.) <https://plato.stanford.edu/archives/win2016/entries/principia-mathematica>.

31. Jha, Saurabh, and Eric J. Topol. "Adapting to artificial intelligence: radiologists and pathologists as information specialists." *JAMA* 316.22 (2016): 2353-2354.
32. Journal Citation Reports. Thomson Reuters.
<https://jcr.incites.thomsonreuters.com/>
33. Knapp, Susan. "Artificial Intelligence: Past, Present, and Future" Dartmouth College. July 24/07. URL:
<http://www.dartmouth.edu/~vox/0607/0724/ai50.html>
34. Laland, Kevin N., and William Hoppitt. "Do animals have culture?." *Evolutionary Anthropology: Issues, News, and Reviews* 12.3 (2003): 150-159.
35. Lindsay, Robert K., et al. "Applications of artificial intelligence for organic chemistry: the DENDRAL project." New York (1980).
36. List of Artificial Intelligence Films. Wikipedia. Accessed: 05-10-07.
https://en.wikipedia.org/wiki/List_of_artificial_intelligence_films
37. "Losing Humanity: The Case Against Killer Robots". Human Rights Watch. Accessed: 05-10-7.
<https://www.hrw.org/report/2012/11/19/losing-humanity/case-against-killer-robots>
38. Machine Intelligence Research Institute (MIRI).
<https://intelligence.org/>

39. McCarthy, John. "A Proposal For The Dartmouth Summer Research Project On Artificial Intelligence" August 31, 1955. <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>
40. McCorduck, Pamela. "Machines who think." (2004).
41. National Motor Vehicle Crash Causation Survey: Report to Congress. July 2008.
<https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/811059>
42. Ng, Andrew. "Artificial Intelligence is the New Electricity." Stanford Graduate School of Business. Feb 2, 2017.
43. Numenta. <https://numenta.com/>
44. OpenCog Foundation. <http://opencog.org/>
45. "Pan-Canadian Artificial Intelligence Strategy Overview". Canadian Institute for ADvanced Research (CIFAR). 30-03-2017.
<https://www.cifar.ca/assets/pan-canadian-artificial-intelligence-strategy-overview/>
46. Patel, Vimla L., et al. "The coming of age of artificial intelligence in medicine." *Artificial intelligence in medicine* 46.1 (2009): 5-17.
47. Price, Derek Solla. "A history of calculating machines." *IEEE Micro* 4.1 (1984): 22-52.
48. Ray Kurzweil, "The Singularity Is Near". Viking. 2005. 978-9-670-0338403

49. Riley-Smith, Ben. "Government to plough £20m into artificial intelligence research including robots and driverless cars" The Telegraph. 26-02-2017.
<http://www.telegraph.co.uk/news/2017/02/26/government-plough-20m-artificial-intelligence-research-including/>
50. Scott, Andrews. Hughes, Barry. "Background Report on Artificial Intelligence." Pardee Center for international Futures
51. Shalal, Andrea. "The Pentagon wants at least \$12 billion to fund AI weapon technology in 2017. Reuters. 12-14-15.
<http://www.businessinsider.com/the-pentagon-wants-at-least-12-billion-to-fund-ai-weapon-technology-in-2017-2015-12>
52. Sullins III, John P. "Gödel's incompleteness theorems and artificial life." *Techné: Research in Philosophy and Technology* 2.3/4 (1997): 185-195.
53. Turing, Alan Mathison. "On computable numbers, with an application to the Entscheidungsproblem." *Proceedings of the London mathematical society* 2.1 (1937): 230-265.
54. Turing, Alan M. "Computing machinery and intelligence." *Mind* 59.236 (1950): 433-460.
55. United States Department of Labor. Bureau of Labor Statistics. "Taxi Drivers and Chauffeurs." Accessed 05-10-07.
<https://www.bls.gov/ooh/transportation-and-material-moving/taxi-drivers-and-chauffeurs.htm>

56. Wajcman, Judy. "Automation: is it really different this time?." *The British Journal of Sociology* 68.1 (2017): 119-127.

57. Web of Science. Thomson Reuters. [https:// www.weboknowledge.com](https://www.weboknowledge.com)