

The Impact and Critical Success Factors of Integrating Artificial Intelligence

MASTER THESIS

by

Devinson Peña

submitted to obtain the degree of

MASTER OF SCIENCE (M.Sc.)

at

TH KÖLN - UNIVERSITY OF APPLIED SCIENCES
INSTITUTE OF INFORMATICS

Course of Studies

WEB SCIENCE

First supervisor: Prof. Dr. Jan Karpe
TH Köln - University of Applied Sciences

Second supervisor: Prof. Dr. Gerhard Hartmann
TH Köln - University of Applied Sciences

Cologne, August 2017

Contact details:

Devinson Pena
Duisburger Strasse 68
40479 Düsseldorf
devpen787@gmail.com

Prof. Dr. Jan Karpe
TH Köln - University of Applied Sciences
Institute of Informatics
Steinmüllerallee 1
51643 Gummersbach
karpe@gm.fh-koeln.de

Prof. Dr. Gerhard Hartmann
TH Köln - University of Applied Sciences
Institute of Informatics
Steinmüllerallee 1
51643 Gummersbach
gerhard.hartmann@th-koeln.de

Abstract

Artificial intelligence (AI) technologies are already today impacting the way organizations are creating strategies, forming their teams and planning their technology investments. With the explosion of available data, fast pace cloud computing developments, continuously improved hardware and a truly global network, in particular machine learning technologies are thriving and forming the backbone of the AI movement.

Use cases across different industries around the globe have revealed in the past years the benefits of a higher degree of automation in daily processes, the speed and accuracy of data analytics performed by machines and increased productivity. By lowering barriers of entry with accessible open source libraries and algorithms, business executives get further triggered to look for the best AI integration plan for their company to stay competitive and to prepare the path for a fast evolving future.

Nonetheless, there is also a vast amount of organizations and decision makers who intentionally or unintentionally haven't dealt with AI yet. Knowledge gaps and ethical concerns are mainly hindering the broad deployment of AI technologies.

In light of this context, this thesis aims at outlining the impact an integration of AI has on a companies strategy, organization and technological investments. It further highlights key success factors and addresses concerns decision makers need to consider when getting started with AI. A recommendation in form of a checklist shall enable decision makers to better understand all decisions which need to be made before integrating AI into their business.

Contents

Foreword	6
1 Introduction	7
2 Thesis Overview	12
2.1 Hypothesis	12
2.2 Research methodology	13
3 Defining Artificial Intelligence	14
3.1 Perceptions of AI	16
3.2 Critical technological observations	20
4 Current AI Capabilities	23
4.1 Current use of AI technology	24
4.2 Key challenges when adopting AI technology	26
4.3 Getting started with AI	27
5 Machine Learning	33
5.1 Machine learning enablers and use cases	34
5.2 Machine learning methods and algorithms	36
5.3 Deep Learning Neural Networks	39
5.4 Implementation of Machine Learning	41
5.5 Considerations for using ML	42
6 Impact and Success Factors of AI and ML	44
6.1 Impact of AI on different industries	45
6.2 Dimensions of impact: Strategy, Organization and Technology	48
6.2.1 Strategic impact	49
6.2.2 Organizational impact	60
6.2.3 Technological impact	65
7 Recommendation	72
8 Further research	77
8.1 Law	77
8.2 Policy and regulation	78
9 Conclusion	80
List of figures	82

Bibliography	82
References	83

Foreword

I would like to thank Linda, Maria and Natalie for all of your support.

1 Introduction

Artificial Intelligence (AI) is not only going to impact how we do business but also will define the type of work we will do in the future. The world has reached a tipping point in our human technological capabilities which now allow AI applications to create value at an exponential rate. There are now low entry barriers in harnessing AI technologies and its applications. Companies are using this moment in time to rethink their current business models, processes and overall strategy to reap the potential benefits of AI. The advantage often goes to the company's most adept at choosing technologies, not to the businesses that create them (Iansiti & West, 1997). This paper will reflect on those businesses and take a deep dive on the impact these businesses will face initiating the process of integrating an AI driven strategy, its technologies, and applications.

Establishing an AI driven strategy will require change. Just as with the digital revolution, AI will need a change in thought, processes, management, strategy, and understanding its importance. Understanding is key to any organization bringing in new technology or tools (Westerman, Bonnet, & McAfee, 2014). Most businesses today have already adapted to the digital age or the "Third Industrial Revolution." Those that have adapted and positioned themselves to capture the real benefits of the digital age have acquired a real "digital advantage", and the majority of companies think that the digital transformation is a competitive opportunity (Westermann, Tannou, Bonnet, Ferraris, & McAfee, 2012). Through this revolution, data generation has been one of its largest byproducts.

Our lives have become data centric, and more and more aspects of our lives and work generate vast amounts of data, which also makes us more predictable. Data is a significant factor in all of this and one of the most critical denominator. It is the creation, access and understanding of this data that's an essential step towards any AI strategy, especially its applicability.

The most common AI application is finding patterns in enormous quantities of data. These applications are reducing cost, improving customer satisfaction, productivity and increase revenues. Many companies are investing billions in integrating AI tech-

nologies and creating applications fitting to their needs. The potential for AI in organizations is enormous. In the coming years the AI market will grow to become a multibillion dollar market and doing business can take up a whole new meaning. As a measure of its significance, according to a McKinsey study, technology giants, Baidu and Google have spent between 20 billion to 30 billion dollar on AI in 2016, with 90 percent of this spent on Research and Development (R&D) and deployment, and 10 percent on acquisitions (Columbus, 2017).

AI technologies and applications today come in the form of robotic vehicles, speech recognition, game playing, logistics planning, machine translation - to mention only a few examples. The term, "Artificial Intelligence" is used loosely and popular among companies trying to pivot with the changes of technology. Truly robust and general AI does not technically exist, yet. This paper focuses on AI referred to as applied, narrow, weak or practical AI and concentrate on the most common technologies for applied AI which most people experience today and companies are integrating, such as Machine Learning (ML). ML is feeding off the adoption of the digital age and creating a gateway to the next industrial revolution or what we call "the fourth industrial revolution".

Two other leading voices in AI have expressed the same sentiment of AI's impact potential. Even Baidu's former Chief Scientist Andrew Ng believes that AI is about to transform many industries, just as electricity did 100 years ago and the internet did 20 years ago. Baidu is a large web services company and one of the biggest internet companies in the world.

As we saw with digital disruption, AI disruption will mean success for some and failure for others. Unlike the internet, those who were late to the game were able to beat all of those who were first to the market; AI is different and companies that get started immediately with AI may enjoy a lasting advantage.

The focus is to provide key decisions makers with an understanding of the advantages and disadvantages to expect when integrating AI into their organizations. Through examples, the challenges most companies will face in this process of integration will be explored. Based on the findings, a recommendation will be derived in form of a checklist of success factors company need to keep in mind when making AI an integral part of their business.

Before decision makers take that step down the AI path, they need to make sure they understand what it is, what it does and what it does not. They will also need to

acquire a comprehensive understanding of its technologies and applications. In the next sections, AI will be defined, the current perception of AI and its meaning for the business will be outlined.

Related Work

The general implications of AI on people and our society is an increasing topic of research, with its resurgence into the mainstream through the promotion of tech giants such as Facebook, Google, Amazon, Microsoft, and IBM. Their direct and indirect influence on their widely used products, platforms, and services that apply AI, Cognitive or Machine Intelligence technology and approaches, are one of the main drivers of investments into AI solutions by businesses and funding for further research.

Research on AI that reflects the impact and changing landscape of businesses which have incorporated AI is currently being conducted by academic institutions, large technology companies, consulting firms as well as individually led projects by enthusiasts and experts. Many are working together to meet this increasing demand for AI. For example, the Boston Consulting Group is sponsoring research and analysis and working alongside with the Massachusetts Institute of Technology on Artificial Intelligence and Business Strategy. The MIT Sloan Management Review has already generated numerous written articles that explore how AI is affecting the development and execution of strategy of organizations.

The majority of the work done so far has provided a great comprehensive and holistic insight into how AI will affect various parts of society and economy. This paper on the other, concentrates specifically on the changes business should expect and actionable steps that can be taken to meet them head on. Let's take a look at the experts and how the work differs.

Behind MIT's research, there are two leading researchers, Erik Brynjolfsson and Andrew McAfee. Erik Brynjolfsson is the director of MIT's Initiative on the Digital Economy, and his research examines the effects of information technologies on business strategy, productivity and performance, digital commerce and intangible assets. Andrew McAfee, the principal research scientist at MIT studies how digital technologies are changing business, the economy, and society. Both have coauthored "The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies and Machine, Platform, Crowd: Harnessing our Digital Future". In the Second Machine Age, they discuss how new digital technologies improve our society while

increasing unemployment and depressing wages. They address society as a whole on how to adapt to these technological advances through funding, updating education and cultivating an interest that supports technology and innovation (Brynjolfsson & McAfee, 2014). In *Machine, Platform, Crowd* they focused on the business implication of these digital technologies. They broke down the oncoming impact of technology in three shifts. The first shift is the move from man to machine through the applications of AI. Machines are taking more control and breaking the standard model of people working with machines by now being more independent such as self-driving cars, fewer cashier in stores, and online language translation (Brynjolfsson & McAfee, 2017). The transition of digital technologies going from only being in the digital world to now making impacts in the physical space is already a disruptive shift. The second shift is from products to platforms, such as Uber, Airbnb, and Amazon. These platforms provide a place for companies and people to create marketplaces which in turn is very valuable in the digital world as they can scale economic benefits for both the users and the owners. The third shift addresses centralized (core) versus decentralized (crowd) institutions which they call “core to the crowd.” The core refers to banks and crowd to self-organized participants such as the formation of Bitcoin, a virtual currency. Due to technology lowering the costs of interacting with experts, more tasks can be achieved outside of the core, which traditionally was kept internally to keep costs lower (Brynjolfsson & McAfee, 2017).

Both provide insight into trends that will push us into the fourth revolution running of the back of the Big Data era. This paper, fills in the gap with insight into how business can start shifting with the trends discussed.

Another significant contributor involved with AI is Tom Davenport. His research and analysis in this area consist of digging into 160 AI projects in which he identified three primary groups of business applications: Robotics and Cognitive Automation (data driven tasks automated through robotics), Cognitive Insights (insights through machine learning detecting patterns) and Cognitive Engagement (interaction with customers and employees). He points out that among these AI applications none of them have wiped out any jobs and highlights that the most successful initiatives involve cross functional teams that combine functionality to perform higher-level tasks (Davenport & Kirby, 2016). Davenport is also the coauthor of “Only Humans Need Apply”, where he addresses the challenges employees will face as machines get smarter and start to take over tasks. Davenport points out that employees will require more extensive experience and education to stay relevant in the future highlighting the estimated 47 percent of jobs computer automation is set to be eliminated in the United States alone. Also,

he stresses the fact that machines would still need human assistance to achieve the best results by combining their processing power with complex human thinking. This contrast of elimination through automation and augmentation of human performance will lead to a market with fewer jobs for an increased number of talented applicants causing a drop in wages (Davenport & Kirby, 2016).

This paper relates closely to the issues addressed by Davenport but instead address the decisions makers who will ultimately make the decision affecting the employees and involvement in the adoption an AI.

Davenport's insights on the impacts of AI are also similar to that of another prominent figure, Jerry Kaplan an futurist, entrepreneur, author and computer scientist. In Kaplan's book, *Humans Need Not Apply* he addresses the future opportunities, challenges, and danger of machine intelligence. Kaplan's work provides an evaluation of what to expect such as the ultra-rich using technology to get even richer even faster. He sees this happening through the advancements in machine learning algorithms that can accumulate profits in productive environments like the stock market through high-frequency trading or companies that can price arbitrage by targeting individual customers (Kaplan, 2015). Also, he warns that just like certain skills becoming obsolete due to automation, it is possible to someday have "artificial persons" managing people.

Kaplan also drives down hard on the fact that machines have the potential to displace human worker or risk being governed by them. This paper takes this into consideration but reflects on the importance on how top management and employees can work together to make sure machines work in tandem as a team and in a race against each other.

Their work provides lots of insight into the oncoming transformational changes AI will have on society and general guidelines on how to take advantage of it while reducing its negative effects. Their work also addresses people that are preparing to enter the workforce or already in the workforce and those that will welcome the next generation of workers. This paper takes these learnings as a foundation and builds on top of them to specifically target the decision makers of today and tomorrow, such as those in the top management who are currently or beginning to address these challenges of transforming their business strategy, technology systems, employee skills, and culture.

2 Thesis Overview

This paper focuses on the impact and critical successful factors of integrating AI with a concentration on ML technology and integration best practices. The paper provides an overview of the changes AI and ML is having on dimensions of an organization that top management have influence over such as the strategy, the people within the organization and technological. Each dimension provides impact and factors that must be considered for successfully integration based on best practices across industries and sectors who have already started the process of integration. From these findings a recommendation is provided in the form of a checklist for top management and managers to use for guidance in an effort to mitigate the friction caused by introducing a new technology into a company, especially one as disruptive as this one.

This paper will be structured as follows:

- **Section 1:** Definition of AI, its perception in general public and business and the technological observations
- **Section 2:** Current AI capabilities and challenges with adoption
- **Section 3:** Machine learning, why it is exploding, and its importance, the expansion of deep learning and how companies are using it
- **Section 4:** Impact and success factors across three dimensions: strategy, organization, technology
- **Section 5:** Recommendation on what companies should consider to successfully integrate AI applications with its people, in its processes, and across its existing technology

2.1 Hypothesis

The underlying hypothesis of this paper is that if top management and important decisions makers understand the impact of AI and ML technologies they will be better informed and increase the likelihood of successfully integrating AI and ML into their

organizations. Success would be based on their ability to create a strategy that fits the business needs, empowers its employees and mitigates technical debt.

If this hypothesis is true, the logical conclusion would be that companies need to acknowledge the impact and address these critical success factors before integrating AI into their current business model or, at the very least, consider them.

If the hypothesis is false, the logical conclusion would be that companies do not need to understand the impact and critical factors of integrating AI and AI could become naturally part of any business.

2.2 Research methodology

This thesis is based on the use of archival data as a method to collect data through documentation that already existed such as organization records, documents produced by individuals, publications, previous research and multimedia materials.

Many sources include as well secondary qualitative and quantitative research conducted by institutions, consulting firms, service providers and enthusiast expert groups in the form of yearly technology reports, surveys, and interviews through webinars and digital podcasts.

Other methods used where historical and grounded theory research based on grounded practices coming from large institutions, associations and technology companies such as the Machine Intelligence Research Institute (MIRI), Association for the Advancement of Artificial Intelligence (AAAI) and Facebook AI Research (FAIR).

3 Defining Artificial Intelligence

Research reveals that there is no all encompassing definition of AI. John McCarthy who created the term “Artificial Intelligence,” defines AI as “the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.” (McCarthy, Minsky, Rochester, & Shannon, 1955).

In “Artificial Intelligence: A modern approach,” Peter Norvig and Stuart R. Russell, broke down AI and categorized it in two dimensions to provide a more comprehensive definition. The first dimension focuses on thinking and acting human and the second dimension on thinking and acting rationally. Rational in this case means doing the right thing given what is known. They define AI as, “the designing and building of intelligent agents that receive percepts from the environment and take action that affects that environment” (Russell & Norvig, 2003).

According to Ben Thompson, AI can be defined in two parts, saying “Artificial General Intelligence (AGI) is a computer capable of doing anything a human can. That is in contrast to Artificial Narrow Intelligence (ANI), in which a computer does what a human can do, but only with narrow bounds.” (Thompson, 2017). It states ANI is capable of accomplishing specific tasks very well yet limited by its combination of data and other machine learning techniques implemented and controlled by people. AGI, on the other hand, would be able to improve itself exponentially. The next step in AGI would be artificial super intelligence (ASI), which means surpassing the human intelligence a billion times more than the smartest person. Despite the various levels of AI, this paper will be focused on narrow AI - and concentrates on the technologies and applications we can either pick, design or teach computers to do for humans based on what is currently possible today.

A common thread among all of these varying definitions is the term intelligence. This term is deemed to a certain extent a philosophical one, which also causes much debate for its definition. A popular one from John McCarthy is, “Intelligence is the compu-

tational part of the ability to achieve goals in the world. Varying kinds and degrees of intelligence occur in people, many animals, and some machines.” (Nilsson, 2010).

The AI100 Stanford study panel uses the calculator analogy to make its point about intelligence. A simple electronic calculator performs calculations much faster than the human brain, and almost never makes a mistake. Is a calculator intelligent? They see intelligence lying on a multi-dimensional spectrum. They go on to explain that the difference between an arithmetic calculator and a human brain is no one of a kind, but of scale, speed, degree of autonomy and generality (Stone et al., 2016). These factors can be used to evaluate every other instance of intelligence. What is stressed by many experts is that AI today, only mimics intelligence in certain ways but it is not intelligence. What we perceive as intelligence is a result of math or correlations of highly-complex cause/effect relationships between inputs and outputs.

The AI100 study panel goes on further to explain that human intelligence is a natural choice for benchmarking the progress of AI. Any activity computers can perform and people once performed should count as an instance of intelligence. Many systems surpass human intelligence in speed, such as scheduling the daily arrivals and departures of thousands of flights in an airport.

Alan Turing in 1950, sought to design a test which would define intelligence through the Turing Test. The goal of the experiment is to assess whether a computer is intelligent or not through a series of questions asked to the computer and a human. If the human interrogator cannot distinguish between the computer and human, then it passes the test (Turing, 1950). Today for a computer to pass the Turing test it would need to account for the following:

- Natural Language Processing
- Knowledge Representation
- Automated Reasoning
- Machine learning
- Computer vision
- Robotics

These are the six disciplines that AI is composed of that would have to function to pass the Turing test (Russell & Norvig, 2003). AI researchers, however, have focused

on understanding the principles of intelligence rather than duplicating the test. Those who have achieved significant progress in the AI, have taken the same approach and use the sciences that govern each area mentioned above to reap their benefits.

As noted ANI benefits generated today are particular in subject matter, and in the task, they are meant to accomplish. Machine learning algorithms and techniques have made this possible in the ANI era. They are built to figure things out by themselves with a key ingredient of data. Data is the “new oil” helping machine learning move forward in long strides. The notion of ANI vs. AGI is clear to those involved in the AI field in some form or fashion but what exactly is the general perception of AI among those who are currently in discussions about its integrations into their business organizations?

3.1 Perceptions of AI

Perception of the general public

Already today, many services are using AI applications that the general public does not recognize. Google’s search engine uses it when it retrieves relevant results; email services use it to filter out spam and e-commerce websites like Amazon use it to recommend products. The list can go on, but one important thing to establish is that the AI technologies and applications today are not close or as entertaining as depicted in their robotic representations from science fiction movies such as in 2001: A Space Odyssey, Terminator, Ex Machina or television series like Westworld.

Unfortunately, films and novels shape a lot of today’s popular imagination and perception of AI. Those with the right perspectives currently are seeing the benefits of AI tech and applications at schools, homes, and hospitals. The R&D departments at universities have understood its potential and have devoted entire departments to AI studies along with technology companies like Apple, Google, IBM and Microsoft who are spending heavily to explore AI. This interest is not new; it has just been reintroduced. Since its founding in 1950, interest, awareness, and discussions exist in AI. Research from Fast & Horvitz found that since 2009, discussions of AI have increased sharply and that the debates around the topic have been consistently more optimistic than pessimistic. They explain in their study that when they examined specific concerns, they found that loss of control of AI, ethical concerns and the negative impact of AI on work had grown (Fast & Horvitz, 2017). Unique to industry though, they found that AI in healthcare and education increased over time. To assure people about

AI, coalitions like the “Partnership on Artificial Intelligence to benefit individuals and society” has been formed by the largest tech companies.

AI allows for innovation to blossom and creates opportunities for both people and businesses, just like the internet led to the boom of large corporations like Google, Amazon, and Baidu to name a select few. However, from a different perspective, the advancement in various AI disciplines is creating concern that these AI technologies may permanently displace human workers, robotic warfare and make Orwellian surveillance techniques easier to develop, among other are its disastrous effects (Markoff, 2014).

From an optimistic standpoint, AI has the potential to augment human capabilities for the better. From a pessimistic perspective, it raises concerns on where we even still need people. It is important that we understand these perceptions and address all related concerns since they can turn into regulatory activity which potentially is a serious repercussion (Stone et al., 2016). For example, there have been suggestions from the government to regulate AI development to prevent threats to humanity (Guardian, 2014).

Fast & Horvitz conducted a study which tracked the impression of AI in the news over the past 30 years, 1986 to 2016, and over 3 million articles from the New York Times. The study itself used crowdsourcing and another AI technology called natural language processing (NLP). Through NLP they extracted mentions of AI and used paid crowdsourcing to annotate these mentions with measures of relevance, levels of pessimism or optimism. Other public opinion polls have similarly measured topics relevant to AI and they support their findings, showing greater levels of optimism than pessimism about AI, but increasing existential fear and worry about jobs (News, 2016).

Perception in businesses

How are companies and the people in them currently dealing with these types of topics within their organizations? This question can be addressed by looking into the perceptions of three key groups in organizations: management, employees and customers.

Management: Most business executives are well aware of the disruption AI is either currently causing or will cause. A survey from the Economist was conducted to gauge the corporate attitudes toward AI across various regions and industries. The global study included 203 senior executives, and more than 40 percent anticipate AI will

start displacing humans in some jobs in the next five years. Slightly more also see AI changing their role in a more augmented way, making their job easier and improving their performance (Economist, 2016).

From a competition perspective, many executives are advancing their efforts in preparation for the disruption they expect AI to create. The survey says that 43 percent of respondents feel that any delay in introducing AI will make them vulnerable and 46 percent worry about AI-based competition from startups (Economist, 2016). According to the survey findings, the executives are expecting AI will create a better user experience, more accurate decision making, and improvement in operating efficiencies and therefore boost revenue.

Before reaching any desired results, executives will need to build comprehensive business cases. From the survey, 29 percent of the respondents said making a business case would be very complicated due to the AI technologies and applications not being mature enough, which correlated to the 30 percent who said senior management lack an understanding of AI, which in turn hinders the business case.

Employees: Job loss is a prime concern for employees. The main question they ask themselves is, “Are my skills still relevant?” and “Will I still have a job?”. There are two opposing views (negative and positive) that touch upon these questions which date back to the 1930’s “technological unemployment” theory addressed by John Maynard Keynes. It suggests that automation negatively affects employment by displacing workers tasks but positively will increase demand for labor in other industries (Petropoulos, 2017). We see this effect at legal firms, where software is capable of analyzing large legal documents, which reduce the cost of research but increased the demand for it. As a result, the number of law clerks rose 1.1 percent on average between 2000 and 2013. We saw the same with the increase in Automated Teller Machines (ATMs) reducing the number of bank clerks but also because of the reductions in costs of running a bank; it allowed it to open more branches in response to customer demand (Petropoulos, 2017). These examples held true for the digital revolution but according to the McKinsey Global Institute, AI’s disruption on society is happening ten times faster and 300 times the scale, meaning 3000 times the impact (Chui, 2015).

PWC conducted a survey called, “The future of work, a journey to 2022”. In that survey 10,000 people in China, India, Germany and the UK and US shared their view on the future of work and what it means for them. The survey found companies had replicated techniques for collecting customer insight for their employees. For example,

sensors to check location, monitor performance and even check employee health is already a reality. Collection of employee data will redefine “contracts” in a way where employees will need to hand over data (e.g. health, performance, possible private life) in return for job security. In the PWC global survey, 30 percent of the participants said they would be happy for their employers to have access to personal data. This 30 percent, were primarily younger adults.

Customers: According to a Weber Shandwick and KRC Research survey of 2,100 consumers online across five global markets and 150 Chief Marketing Officers (CMOs) responsible for marketing and branding, they found consumers are accepting of AI. Consumers saw AI having a positive impact rather than a negative one (52 percent and 7 percent). It is important to note that about 18 percent said they knew a lot about AI and one-third acknowledged knowing nothing. Despite not knowing anything about AI, the most surprising was the results on trust into the results of AI (Gaines-Ross, 2017). Two-thirds or more stated that they would trust in AI used for medication reminders, travel directions, entertainment, targeted news, and manual labor and mechanics. Out of all the tasks listed in the survey, child care ranked at the bottom of the list - a task where customers seem to have their doubts that AI is able to replace. Consumers are also employees at other companies and they too, are worried about staying relevant, job loss, data security, and privacy. From the survey, 82 percent said AI would lead to more job loss than creation (Gaines-Ross, 2017). Half indicated their concern for cyber-attacks and stolen data as well. In general, the findings of the survey concluded that the consumers are ready to accept what they think AI is at the moment.

In conclusion, AI is perceived differently in public as well as in businesses depending on the individual experience with AI. People acknowledge its benefits once they are aware of it, but a lack in understanding and a lack of direct interaction with AI for some, also give room for fear and insecurities. There is no doubt, that AI is changing how people interact. People’s relationships with machines will become more prominent as AI systems learn to adapt to people’s personalities and goals in all types of settings - be it at the workplace or privately at home. The two standpoints regarding AI - the chances it brings along on the one hand and the risks people fear on the other hand - are also discussed controversially in the scientific world. For example, two renowned experts in their respective science fields have vastly distinctive public opinions. On the one hand, we have Stephen Hawking warning ASI, “could spell the end of the human race,” while Ray Kurzweil believes it would cause a “technological singularity” which would lead us to some shape or form of utopia.

The next chapter should help to understand why experts think that the time for AI is now.

3.2 Critical technological observations

There are individual views from various visionaries who have given us insight into the progression and path technology has taken and will take. When it comes to AI, Ray Kurzweil, a computer scientist, futurist and currently hired by Google to work on machine learning and language processing projects, has been the most popular of the visionaries. His observations provide insight into how we have managed to get where we are today with AI technologies and applications. These observations have aided many in what to expect from our technological future and prepare for it, which is why anyone interested in integrating AI should be aware of them.

The law of accelerating returns

The ability for us humans to progress quicker is what Ray Kurzweil calls “Humans history’s law of accelerating returns.” The law points out that advanced societies have the ability to progress at a faster rate than less advanced societies.

The 19th-century and 20th-century humanity knew more and had better technology than 15th-century humanity, so it is no surprise that we currently see far more advancements than the previous generations. In general, an analysis of the history of technology shows that technological change is exponential. Kurzweil states, we will not experience 100 years of progress in the 21st century, it will be more like 20,000 years of progress at today’s rate (Kurzweil, 2001).

Kurzweil speaks of technological change in particular paradigms, methods or approaches to solving a problem. He states that paradigms provide exponential growth until the paradigm exhausts its potential. At this point a paradigm shift occurs in the technological approach, enabling the exponential growth to continue.

Kurzweil explains, “Moore’s law of integrated circuits was not the first but the fifth paradigm to provide accelerating price-performance. Computing devices have been consistently multiplying in power (per unit of time)” (Kurzweil, 2001). The five paradigm shifts have been:

1. Electromechanical computer - built by IBM for the 1890 U.S Census

2. Relay based computer - Alan Turing's relay based computer cracked the Nazi Enigma
3. Vacuum Tube - the vacuum tube computer predicted Eisenhower's win in 1952
4. Transistor - transistor based machines were used in the first space launches
5. Integrated circuit based personal computer

Kurzweil expects the 6th paradigm shift to be in the three-dimensional molecular computing and the hardware to give rise to human level intelligence in machines (Dorrier, 2016).

The power of exponential growth

Another observation by Kurzweil is the power of exponential growth. According to him, we as people still underestimate the progress that is coming because it is hard for humans as naturally linear thinking creatures to internalize exponential technological change.

It is important to get an understanding of the massive scale of advancements that the technologies of the future will enable. Especially now, as we have reached what Kurzweil calls the "second half of the chessboard." (Kurzweil, 2001). The second half of the chessboard is the favorite example used and a valuable lesson for those who are thinking of investing in AI. Kurzweil likes to use this folktale, to drive home the power of exponential growth. This example, Kurzweil talks about an inventor that created chess, a game that demonstrated the importance of all members of society (e.g. king, queen, knight, bishop, castle, and pawn). The emperor at the time was impressed and offered a reward. The inventor asked for one grain on the first square of the board and that it should double on each subsequent square. On the 32nd square, the amount of rice amounted to a production of a few acres. By the 64th square, the estimated total amount of rice was comparable to a pile the size of Mount Everest. There are various versions of how this ends. In one version of the story, the emperor is going bankrupt as the 63 doublings ultimately totaled 18 million trillion grains of rice.

The first half of the chessboard for emperor and inventor was uneventful, but as soon as the inventor had accumulated much more than previously thought, the Emperor took notice. Today, technology like a smartphone is a great example to realize how quickly things are accelerating.

In summary, these two technologically relevant observations demonstrate not only the impact that AI can potentially have, but also reveal that the fast pace of technological change will continue allowing AI to spread and grow its reach to the maximum in the upcoming years. The following section will give insights on the types of AI that are currently available and applicable for businesses and will outline ways to pick the most suitable approach for the company.

4 Current AI Capabilities

AI technologies and applications have made many promises for decades and only recently are we seeing some significant results. There are many fields of AI, some of the most prominent ones are robotics, computer vision, speech recognition, natural language processing (NLP) and machine learning (ML). Out of all fields, ML is achieving the most traction in adoption. According to a company AI investment report, ML attracted nearly 60 percent of that investment of all AI (Columbus, 2017).

ML has become far more accessible than ever before outside of its regular expert group of data scientists, data miners and software developers. There have been many barriers to entry in ML, one primary one being the lack of having sufficient data. However, with large amounts of data being generated and stored by companies, ML has been given the opportunity to expand significantly. Due to this fact, there will be one section dedicated to ML in this paper, detailing out what its capabilities are.

There have been some misconception of AI that must be addressed and understood before going into detail. One of the largest misconceptions is that simply saying business is “AI-powered,” doesn’t necessarily mean the company is much more efficient or productive. AI capabilities are not primarily integrated into companies to replace humans but to leverage the chance to augment humans in the organization. Both, machines and humans have their own advantages - where machines can be quicker, cheaper and more consistent than humans, humans have the ability to perform more complex tasks and apply their knowledge broadly. A blended approach between humans and machines is what AI experts strive for.

Another fallacy is that of the power of the algorithm. AI and algorithms are not synonymous. Algorithms are a necessary component of running AI powered technology, but they do not define them. Many experts claim data, not algorithms are the key limiting factor to the development of human - level artificial intelligence. This claim came from Alexander Wissner-Gross who reviewed the timing of most publicized AI advances over the past 30 years and found evidence to suggest that major AI breakthroughs have been constrained by the availability of high-quality training data sets

as seen below (Wissner Gross, 2016).

Year	Breakthroughs in AI	Datasets (First Available)	Algorithms (First Proposed)
1994	Human-level spontaneous speech recognition	Spoken Wall Street Journal articles and other texts (1991)	Hidden Markov Model (1984)
1997	IBM Deep Blue defeated Garry Kasparov	700,000 Grandmaster chess games, aka "The Extended Book" (1991)	Negascout planning algorithm (1983)
2005	Google's Arabic- and Chinese-to-English translation	1.8 trillion tokens from Google Web and News pages (collected in 2005)	Statistical machine translation algorithm (1988)
2011	IBM Watson became the world Jeopardy! champion	8.6 million documents from Wikipedia, Wiktionary, Wikiquote, and Project Gutenberg (updated in 2010)	Mixture-of-Experts algorithm (1991)
2014	Google's GoogLeNet object classification at near-human performance	ImageNet corpus of 1.5 million labeled images and 1,000 object categories (2010)	Convolution neural network algorithm (1989)
2015	Google's Deepmind achieved human parity in playing 29 Atari games by learning general control from video	Arcade Learning Environment dataset of over 50 Atari games (2013)	Q-learning algorithm (1992)
Average No. of Years to Breakthrough:		3 years	18 years

Figure 4.1: Datasets Over Algorithms

Source: Datasets Over Algorithms. Edge (2016)

The following section explores the latest advances in AI technologies, the different options companies currently have in choosing the most suitable AI technologies and applications that are available to them.

4.1 Current use of AI technology

The AI business has exploded, with 5 billion dollar in venture investment, a few big acquisitions and a substantial increase in people reading up on AI developments. The impact of AI technologies and applications is here. Almost every industry is already being affected. Every employee could use AI to become more productive with tools that exist today. Companies have at their disposal for the first time all the building blocks to begin embedding AI technology in their businesses.

The latest advances in AI have been in two broad areas: perception and cognition.

In perception, the most common use and advances have been in the speech and voice recognition area. Millions of people use voice recognition technology in their daily lives using Siri (Apple), Alexa (Amazon) and Google's Assistant. A study by James Landay and his colleagues found speech recognition is about three times faster on average with a 4.9 percent error rate than typing into keyboards on cell phones (Bjorn, 2016). In the same area of perception, another technology that has improved vastly is image recognition. Facebook for example and other apps now have the ability to recognize

friends faces. Other apps help recognize birds and even at corporate headquarters things like ID cards are disappearing and being replaced by face recognition. Apple, for example back in 2013 filed a patent for facial recognition to unlock their iPhone devices (Crook, 2017) . With self-driving cars, vision systems have improved from making a mistake every 30 frames in identifying pedestrians to once every 30 million frames. This improvement has been accelerated by new ML approaches like “deep learning” through neural nets.

The other type of major improvements has been in cognition and problem-solving. Recent popular examples of machines achievements in this area have been shown through Google’s DeepMind team beating a human in poker and the game of Go through the use of ML systems. Other examples are seen in cybersecurity through the use of intelligent agents that can detect malware created by a cybersecurity company “Deep instinct” or Paypal who use it to prevent money laundering. ML is helping Wall Street companies to execute trades and Amazon to optimize inventory and improve product recommendations to customers. ML, in general, is replacing older algorithms in applications, and is superior at many tasks previously done by humans with error rates of 5 percent; for example, the ImageNet database (an image database) is better compared to humans. This error rate threshold created many possibilities for transforming the workplace and economy, and as ML applications surpass human performance at a task, they are much likelier to spread quickly. For example, companies like Aptonomy and Sanbot are working on improving vision systems to automate the work of security guards. Other companies like Affectiva are creating software to help recognize emotions such as joy, surprise, and anger or even use medical images to help diagnose cancer like Enlitic.

The applicability of AI is still narrow when utilized in a noisier environment filled with unknown data or assumptions that are not accounted for by the ML systems. ML systems today are trained to specific tasks and the knowledge is not generalized. The perception that ML technology applies its knowledge to a broader understanding is one of the biggest confusion in the progress of AI. This is discussed in more detail later in the paper.

As the technology stack for AI matures, more companies will increase their capabilities, and new ventures will use AI to fill in gaps where they see the most opportunity. For the past three years, there has been an attempt by O’Reilly Media and Bloomberg to provide an overview of what AI is or as they call it “machine intelligence” landscape looks like. This consolidation of established companies and startups provides a glimpse into how big the AI business currently is, and we are still only at the beginning.

This landscape is very useful for companies looking for vendors already servicing their similar needs to their current business and processes. Although AI is already in use by many companies, most of its opportunities are yet to be tapped into (Brynjolfsson & McAfee, 2017). Despite the vast selection of currently available vendors with expertise, adopting AI is not a plug and chug play activity.

4.2 Key challenges when adopting AI technology

Despite its benefits, adopting an AI technology comes with some challenges. These challenges will depend on having a clear strategy, talent, and infrastructure for its application. These challenges are not industry-specific but generally affect the ability to adopt AI technology. According to Forrester's TechRadar report on AI Technologies, the following challenges are currently being faced by the business world (Gownder, 2017):

- **Having a clear business case:** The ability to apply AI that meets specific business objectives with tangible results is still very challenging. Companies need to create clear investment plans that reflect a Return on Investment.
- **Investing long-term time and effort:** The effort and time required to get AI systems up and running depends on the use case and other variables such as the useable data and skills to work with them. Lots of the machine learning algorithms are available, programming them is not the biggest issue, but rather it is tweaking and training the system with well-curated data. For example, IBM Watson's new diagnostic domain takes between 6-18 months to setup depending on the complications and the available data. Expectations need to be realistic about what's available, and the tradeoffs that need to be considered to achieve speedy integrations when accuracy is not the primary concern.
- **Having AI talent available:** The ability to just apply AI software to data purchased through a vendor or downloaded from an open-source site is not possible. Needed expertise depends on the use cases and the AI techniques. The AI talent pool of specialist is already limited and for businesses, even more so. Even if available, most specialist or researchers are in academia, and only a few can translate their skills over to a business context. AI talent is essential but difficult to find. It is important to assess the needed skill set to adapt software to specific use cases and prepare the data accordingly to maintain the technologies and applications.

- **Providing a robust data management platform:** Around 43 percent of analytic professionals say, "ensuring data quality from a variety of sources" is their biggest challenge (Gownder, 2017). The lack of quality data makes it very hard for companies to deploy AI since it needs large amounts of training data to learn and perform. AI will only be as good as the data available, and it will require strong management and governance.
- **Establishing change management processes and practices:** Since the impact on the organization itself can be significant, establishing a change management process is the biggest challenge to adopting AI technologies. In a Forrester survey, one-third of respondents found change management as one of their most significant risks.

More details regarding these challenges, impacts, and ways to overcome them will be discussed in chapter of this paper.

4.3 Getting started with AI

When it comes to developing AI capabilities, very few companies are relying on internal assets. Many firms look toward crowd sourcing and other forms of open innovation to develop AI applications. Many companies are either partnering with other organizations inside and outside of their sector or acquiring or funding technology start-ups. Partnerships, in particular, are allowing companies to sidestep silos to implement new processes or services based on new technology.

Source: Economist. (2016)

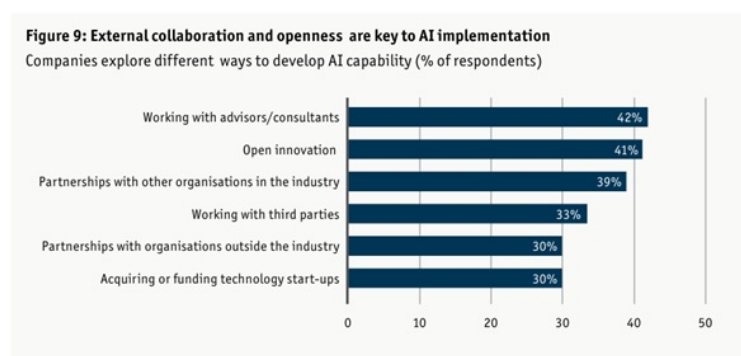


Figure 4.2: Companies explore different ways to Develop AI Capabilities

According to a survey from the Economist, 42 percent of companies are looking to work with advisors and consultants, and 41 percent are open to open-source innovation. On the other end, only 30 percent indicated that they are working on acquiring or funding start-ups or partnering with organizations outside their industry as seen below. The following reflects Davenport's approaches.

So the question arises, how organizations pick the right strategy?

The Resource Pathway Framework is an approach used by most firms looking for growth strategies either by organic (build), partnerships (borrow) or buying the capabilities or the company that offers these capabilities (buy) (Capron & Mitchell, 2013). The Build, Borrow and Buy approach or "The Resource Pathway Framework" was based on Laurence Capron and Will Mitchell decades of research which reflects on the best practices from a variety of successful global organizations path towards growth. This framework is apparent in Tom Davenport's approach introduction of AI into an organization. Davenport narrows down the approaches to three broad categories that he thought were most common:

1. Mostly buy
2. Some buy, some build
3. Mostly build

Each of these approaches have then a further split of different ways to tackle them (Davenport & Kirby, 2016).

The different approaches towards getting started with AI depend on mainly size of the organization and the nature of the products or services the company is offering.

The "mostly buy" approach, for example, is seen in enterprise software companies with AI capabilities built into their product offerings. These companies are either offering AI capabilities that span their whole product set or businesses that specialize in a particular domain. These companies create products for the workplace and serve many markets. These companies can be clustered as follows.

- **Large enterprise vendors:** These larger companies typically have Software-as-a-Service (SaaS) business models like Microsoft suite and Salesforce's Customer Relationship Management Platform (CRM) applications (e.g. Sales, marketing). They use a combination of data science, ML, and deep learning. For example,

Microsoft announced it would make available a free add-on that translates presentations into real-time for their Powerpoint product. This add-on will allow audiences to select their desired languages in the form of subtitles (Znet, 2017). Another example is Salesforce, who built an AI technology called Einstein. Einstein can customize analytics and insights and can sort test by negative, positive or neutral sentiment. They also just recently announced an Einstein Object Detection, which will help with detecting location, size, and quantities of objects, an ideal capability for counting inventory. Some of these companies are even working together, like Salesforce and IBM's Watson AI business. Watson services and data spans various industries like commerce, education, financial services, and health. With this partnership, Watson Insights will now be accessible by Salesforce and mix with its customer relationship data to help better target customers.

The benefit of these vendors is that it does not require AI experts to work with their software. Most of them are cloud-based solutions that charge on a per user subscription basis. Their packages typically come with an integration fee, support teams, and training to get employees started. The disadvantage is that if a company goes all in, it can only go as far as the vendor's platform permits or advance as fast as the vendor does.

- **Specialized vendor:** The specialized companies are ideal for smaller projects that could benefit from AI applications in a particular area of business. Cognitive Scale is one of these companies. They specialize in building cognitive applications in 10 hours, customizes it in 10 days and goes live in 10 weeks, for specific tasks. For one of their clients, M.D Anderson Cancer Center, they build a way for patient families to find lodging and dining while charging it to the patient bills. There are other types of these companies that are more specialized in particular functions and provide AI software solutions. For example, there are Business Intelligence vendors like Narrative Science and Yseop who help automate reports and generating narratives using Natural Language Generation (NLG).

The benefit of these type of companies is that they do not cause significant organizational impacts. They either build the software for the client or provide consulting and training to configure the software for a specific task. The disadvantage here is, similar to the large enterprise vendor, is that there is no control of the product roadmap and how it is built to manage things like quality, accuracy, and speed of the solution.

The second approach is the “some buy, some build” approach. This method provides the ability to have more control. However, this also means it will need talent that can

manage not only choosing the right AI technology but also building and connecting it into the business requirements and goals. There are three ways to implement this approach:

- **Extend capabilities with machine learning:** Machine learning is a method of analyzing data (old or new) and automating the model built. Current techniques of machine learning such as regression analysis are helping to extend a company's analytical capabilities. With the growing amount of data in volume and variety, machine learning is helping produce models that can analyze large, complex datasets faster and build models with more accurate results. There are many models a company can generate using various techniques and depending on the software they can use a vendor or open-source; it can either be very straightforward or complex. This method requires expertise with data science. Machine learning is one of the leading AI technologies and applications integrated into companies. A deep dive into ML will be provided in the next chapter.
- **All-inclusive consultative services:** Large vendors like IBM and Microsoft have built full "cognitive services," which medium to large firms can all take advantage of to introduce large-scale cognitive technologies into their businesses. IBM, for example, likes to initiate their services with a "cognitive value assessment" to assess the best way to use their services. After this, they provide consultants and experts in the field to help implement their solutions. This approach is costly and time consuming but many companies do trust companies like IBM and Microsoft, and would feel comfortable using them as a way to get a start on their path to becoming prepared for AI.
- **Plug into 3rd party API's:** There are many application programming interfaces (API's) that allow programs on your website or app to communicate with medium-level AI technology. These API's provide the opportunity for companies to enhance their services or create additional ones. For example, many companies are connecting to natural language conversation API's which are enabling them to create Chatbots. Chatbots are computer programs designed to mimic conversations with human users. Most chatbots today are built with rules-based decision trees. They are now incorporating NLP, which extracts the user's intents out of sentences and turns it into a command the program can take action. Companies like Google (API.AI), Apple (Siri), Microsoft (Cortana) and Facebook (Messenger Platform) all have platforms for businesses and their developers to build and launch their chatbots.

The third choice a company can make is to create their AI technology stack on its own. In this scenario, a company can use all the tools mentioned previously with the freedom to choose where it is most applicable to their business, with their people or a mix of outside services. This option places all the risk and responsibility on the company and its people to create value through AI applications. When building own AI capabilities, the two focus areas below could be part of the approach:

- **Smarter and more autonomous decisions:** In machine learning, we use various techniques to help build models that help predict with high accuracy the desired output. This output may come in different forms, for example an Alert, specific KPI or a custom metric. In most cases, it is ultimately used to help people make informed decisions. Companies could develop AI technology with more autonomy to create higher levels of productivity. In essence, this would mean letting the output of a model make the decision or trigger an action on its own, with little to no human supervision. This trigger or action of autonomy can come in the form of answering a person's question to a robot or drone moving through its environment. This sort of work demands lots of careful planning and expertise in AI technologies and applications. More physical representations of autonomy are now mostly associated with automobiles and robots. Other than the physical descriptions, we can also find more and more freedom in digital services. For example, in digital advertising, many of the advertisements shown to a user through website banners, videos or mobile apps are delivered autonomously if the user matches the advertiser's criteria. Humans supervise these delivery platforms, but the majority of them make their own decisions.
- **Open-source software:** All of the main technology companies today like Microsoft (CNTK), Google (Tensorflow), Facebook (Caffe 2) and Amazon (AML) have released open-source algorithm software libraries for machine learning and deep learning. This software is free and typically require data scientists, which are hard to find and are expensive people to hire. It also means making a long-term investment to build out specialized capabilities from scratch. In general, open-source software is more secured. It benefits from having large communities having access to it and providing feedback on how to improve or fix it. By using open-source, companies do not have to worry about such things like licensing or subscriptions per user.

In conclusion, there are many different approaches and ways to integrate AI technologies and capabilities. There is no right or wrong combination of methods. Companies can choose one or combine multiple of options depending on their needs. Each combination will demand and require four essential ingredients: data, computing resources

(i.e. hardware), algorithms (i.e. software) and talent to put it all together (Cronin, 2017). The next chapter explores how ML uses all these ingredients to create value for the companies who have deployed it into their current business.

5 Machine Learning

ML and AI today are used interchangeably by big brands and companies that want to announce their latest innovation. However, ML and AI are quite two distinct areas of computing. Machine learning is a subfield of AI. ML like AI also has various definitions but commonly defined across experts in the following way. Arthur Samuel in the 1950's defined machine learning as a field of study that "gives computers the ability to learn without being explicitly programmed" (Samuel, 1959). In more recent days, Tom Mitchell's definitions of machine learning have been more commonly referenced, "A computer program is said to learn from experience E concerning some class of tasks in T , as measured by P , improves with experience E " (Mitchell, 1997).

The common point between both definitions is that machine learning enables a machine to learn. Just like humans interact with the world to understand it, computers can now interact with the world through unique mechanisms that transform its interactions into data. The ability for machine learning programs to learn from experience versus being explicitly programmed for a particular outcome is fundamentally different which in effect creates great value. No longer do we need to code existing knowledge and procedures into machines, which was a fundamental weakness. For example, it is nearly impossible to write down a code in the form of instructions that enable another person to know how to ride a bike or recognize a person's face, although we are aware of how to do it ourselves. Our ability to know more than we can express in code is known as the Polanyi's Paradox from the philosopher and polymath Michael Polanyi which he described in 1964 (David, A. 2014). This paradox has been a key limitation to our ability to code and push machines forward until now with the machine and deep learning; where machines don't need code but examples using structured feedback to learn and one day solve problems such as driving a car.

ML is used and interacted with more than we know. For example, ML is active when Google search searches or when Amazon recommends a product. ML will be the next big wave of innovations to come (Domnigos, 2012). Deep learning (DL) is a subfield of ML enabling this next big wave of innovation. DL is influenced by how the human neurons work. DL focuses on creating artificial neural networks that

imitate the interaction of neurons in the human brain. Facebook has been using it to build systems capable of answering questions not seen in the past, while Microsoft is using it for generating instant translation from English to Mandarin, and many more companies are deploying DL to improve their product and services, such as Google. Google deployed DL on their Android voice search, and the errors dropped by 25 percent overnight and were behind the defeat of the best players of Go, the world's most complex board game.

5.1 Machine learning enablers and use cases

With an increase in data, more powerful computer hardware and the promises of deep learning, machine learning is seeing more breakthroughs than it did since its introduction in 1950's. Here are some important factors that are enabling this adoption:

- **Data:** Data signals from smartphones to industrial equipment, social media, and many other sources combined have created an era of data abundance. In the past two years alone over 90 percent of the digital data in the world has been created, and with the booming internet of things (IoT) connecting billions of new devices and streaming data, we will have much more digital data to work with within the coming decade. Data has become a differentiating and defensible resource for many companies. For example, companies like Google and Facebook invest billions to keep providing services that collect information about their users and their behaviors. The cost of storage of this data has dramatically dropped. For example, Google Compute Engine is seeing a 32 percent reduction in prices and their cloud storage price has dropped 68 percent to just \$0.026/month per gigabyte and \$0.2/month per gigabyte/DRA (Lardinois, 2014). More storage means more businesses can keep lots more data without worrying about space. Data makes algorithms more efficient, and it improves the more data it receives.
- **Hardware:** The hardware used to process data and run these algorithms has been improving continuously. Moore's law, which stated integrated circuits would double every 18 to 24 months, is now celebrating over 50 years is still going strong but running up against the limits of physics and many say it will slow down or level off (Denning & Lewis, 2016). In the meantime, there has been a trending shift from Central Processing Unit (CPU) to Graphic Processing Units (GPUs). GPU's are turning out to be very efficient. GPU's were initially developed to display graphics in computer games, but have been found to be very useful when used for deep neural nets. Neural nets speed up when moved from CPU's to GPU's by 10 times (Buck, 2015). The speed is due to GPU's being able to

process events in parallel. The current market leaders of this type of hardware are Nvidia, Google, and Microsoft.

- **Globalization:** Global networking is also a big enabler of the machine learning movement. With billions of people connected to mobile devices, such as smartphones it is now easier than ever to deliver AI technologies to everyone anywhere on the planet (Jia & Vajda, 2016). For example, the intelligent assistants used in our smartphones are AI. Cloud-based AI is also a big accelerator of learning. For example, if a robot that uses a knowledge-representation system is compatible with other machines, it can upload what it has learned into the cloud and share it with other robots (Brynjolfsson & McAfee, 2017). In this way, robots can gather data from hundreds, thousands or eventually millions of eyes and ears by combining their information into a single system that can distribute the knowledge vastly and rapidly.

The reason for ML being the most popular AI technology was examined by Nils Nilsson with the result (Nilsson, 2010) that ML performs tasks with vast amounts of data which are a lot harder for a human to do ML. The tasks are:

- **Identifies relationships:** Machine learning can identify relationships and correlations in large piles of data where it is complicated for humans to do so in a short period or in cases where the person cannot define a definite relationship between input and output datasets.
- **Fills in knowledge gaps or improves on unknowns:** Humans often don't capture all the nuances in specific environments when designing a program either due to many unknowns or too much information for the human to encode. With machine learning, people can build more holistically and then use machine learning methods to improve itself based on what it learns from its environment, thus making up for the possible gaps missed. For example: driving a car, identifying faces in pictures, naming objects in pictures or even predicting customer churn.
- **Adapts to changing environments:** Many environments are not constant. With machine learning, it is possible to adapt and change with environments especially with unsupervised learning methods. This capability reduces the need for a human to redesign programs to interpret the change and what it means.
- **Analyzes and applies new data quicker:** New information generates or is discovered regularly and sometimes faster than humans can digest. Machine learning helps to keep track of this information and to analyze its significance

more efficiently and quicker than a human would through traditional statistical practice and inference.

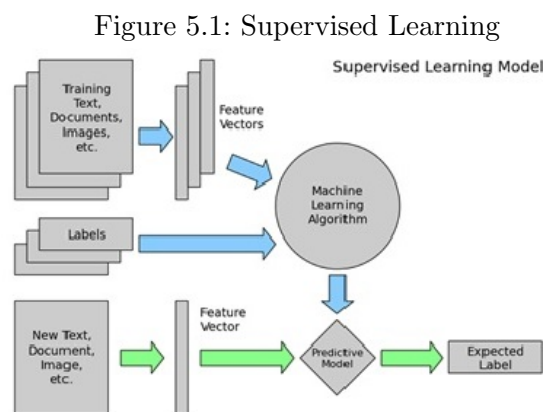
Overall, ML enhances the human's capability to solve problems and make informed inferences by its ability to analyze or find patterns in the major complex data sets faster and with high accuracy. Therefore, capacity to process, analyze and provide insight into large sets of data is critical. Without data, there is no learning done, but with lots of data, we have a lot to learn.

5.2 Machine learning methods and algorithms

There are different types of learning approaches, techniques, and algorithms to choose from within the realm of ML. The most common machine learning approaches used includes supervised learning, unsupervised learning, and reinforcement learning.

The supervised learning approach and the algorithms used under this method are conceptually easier to understand for those starting out in ML. At its simplest level, supervised learning trains model with labeled input and output. The supervised learning algorithms use labeled input and output data to find relationships for patterns and make predictions. For example, inputs might be a picture of various animals and outputs is the right animal labeled in the pictures. After training the systems, it can then be put into production and used in new examples, and if well trained, it can predict answers with high accuracy.

Source: <http://scikit-learn.org/>



Here are some examples of what supervised learning through machine learning can do based on a Tom Mitchell and Michael I. Jordan's progress in the field (Brynjolfsson &

McAfee, 2017):

Input X	Output Y	Application
Voice Recording	Transcript	Speech recognition
Historical market data	Future market data	Trading bots
Photograph	Caption	Image tagging
Drug chemical properties	Treatment efficacy	Pharma R and D
Store Transaction details	Is the transaction fraudulent?	Fraud detection
Recipe ingredients	Customer reviews	Food recommendations
Purchase histories	Future purchase behavior	Customer retention
Car locations and speed	Traffic flow	Traffic lights
Faces	Names	Face recognition

The common supervised algorithms that help predict and train ML models are regression, anomaly detection, and classification algorithms. Below is an example of the types of ways businesses can use the supervised learning approach and algorithms:

- Predicting values (regression algorithms):** If a company wants to predict sales figures, product demand, prices or in general predict a continuous set of variables, they use regression algorithms. For example, Descartes labs uses machine learning to figure out how healthy the corn crop is from space. Corn yield prediction is a big business in the US and being able to predict where the corn market will be in and how it will behave is a significant trend (Brokaw, 2016). The United States Department of Agriculture (USDA) deploys hundreds of workers a year to survey thousands of farms a month. Descartes lab employees are only 20 people and they never leave the office from Los Alamos, New Mexico. Since 2014, their corn yield estimates report has have out-predicted the USDA on all points beating USDA's accuracy by a percentage point. Their algorithms, when compared to historical backtests, show that they have gotten their margin of error down to a 2.5 average.
- Find unusual occurrences (anomaly detection algorithms):** When there is a need to find unusual events or what is called "anomaly detection", such as credit risk, detect fraud or capture abnormal equipment reading, there are specific algorithms used. Paypal has been working with machine learning based pattern recognition for several years, specifically using a mix of supervised and unsupervised deep learning to help detect complex patterns of cybercrime and online fraud. The algorithms analyze tens of thousands of possible features such

as time signals, actors and geographic location that pertain to a type of fraud, subtypes or variants of the same scheme. Once an ML model detects a potential fraud case, humans "detectives" assess what's real and what's not and what to do next (Harris, 2015).

- **Predict categories (classification algorithms):** The most common of all scenarios are predicting types. The difference between predicting categories and predicting values is that predicting categories has a finite set of classes to predict as opposed to predicting continuous variables which can generate an infinite amount of values. Predicting categories can be simple or very complex. To answer questions that are simple such as yes or no, true or false or a set of two different questions use a two-class classification algorithm. A multi-class classification algorithm is used to solve complex tasks categorized into 4-5 answers such as moods or maybe an image resembling a particular figure with over 20 different kinds of possibilities that need 20 different individual classes.

Unsupervised learning is distinct to supervised learning in that the algorithms learn by example, but the examples are not labeled. So for example, if we wanted to train a supervised learning algorithm to identify cats we would show it pictures of a cat and without a cat, labeling the ones which contained a cat. Unsupervised learning uses artificial neural networks (explained further in the paper), which makes it possible to skip the labeling and only present it with pictures of a cat (Kaplan, 2015).

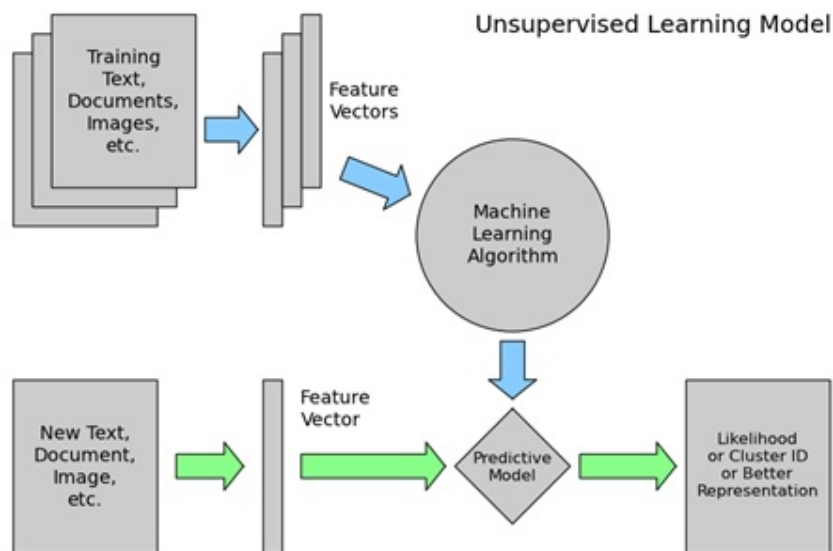


Figure 5.2: Unsupervised Learning

Source: <http://scikit-learn.org/>

We as people can be categorized for example as unsupervised learners since we learn on our own with little to no labeled data. It is harder to learn this way for machines but is very promising through machines who can sift through millions of data points. Unsupervised learning can identify finer correlations between data points within large data sets which otherwise would be too hard for humans to identify. Unsupervised learning clusters inputs and helps to put new information into an appropriate clusters. Hence the use of clustering algorithms is ideal for discovering structure using the unsupervised approach.

- **Discover structure (clustering algorithms):** When there is a need to discover structure such as perform customer segmentation, predict custom tastes or determine market prices we use clustering algorithms such as k-means, patient and related algorithms that can carry out this task. With clustering algorithms, businesses can confirm assumptions about types of groups or identify unknown groups in their complex data sets.

Another learning style is reinforcement learning. It is a form of learning that helps maximized toward a goal in a particular context based on rewards. Accomplished by the machine taking action and then comparing the outcome to a defined reward, it will try to figure out how to maximize the rewards by itself. Today we see reinforcement learning applied in various places like manufacturing with robots that need to pick up and put objects into specific locations or in the finance sector for trading stocks.

Many machine learning tool sets provided today have cheat sheets and extra documentation to provide more guidance on what algorithms to use best. In most cases, these tool sets are setup to apply algorithms to the data set which can help train models. In summary, ML is about looking forward and predicting values, detecting values, classifying values or identifying unknown values within various scenarios. In business, it is important to be pragmatic and realistic. ML creates value within a business process when it is repeatable, reliable and executable (Dunning, 2017).

5.3 Deep Learning Neural Networks

It is important to understand the disruptiveness of deep learning neural networks (DNN). As with traditional machine learning algorithms, we might not be aware of all the examples of deep learning around us such as:

- **Google Translate** - uses deep learning and image recognition to recognize written and spoken languages.

- **CamFind** - uses deep learning to take pictures of objects and tell you what they are.
- **Digital Assistants** - uses deep learning to process natural languages (NLP, extracts meaning) and recognize speech (ASR, determines words based on acoustic signal).
- **Recommendation Engines (Amazon, Netflix, Spotify)** - uses deep learning to recommend next best offers, movies or music
- **Paypal** -uses deep learning to prevent fraud in payments

DNN is a type of algorithm that uses artificial neural networks (ANN) to create layers of representation and abstraction that helps make sense of data such as images, sound, and text. The uniqueness about DNN is its ability to extract features (also referred to as neurons) automatically that are relevant to the solution of the given problem that was trying to be solved such as knowing if an image is a dog or not without labels (as referred to in unsupervised learning approach). DNN can be used to augment supervised learning and unsupervised learning problems. It essentially helps to reduce the burden on a programmer labeling features.

Each DNN layer is responsible for training a unique set of features based on the output of the previous tier. For example, as the DNN algorithm gets fed more pictures of a dog, the DNN layers increase, and it forms a hierarchy of low-level features to high-level features based on its ability to identify the dog accurately.

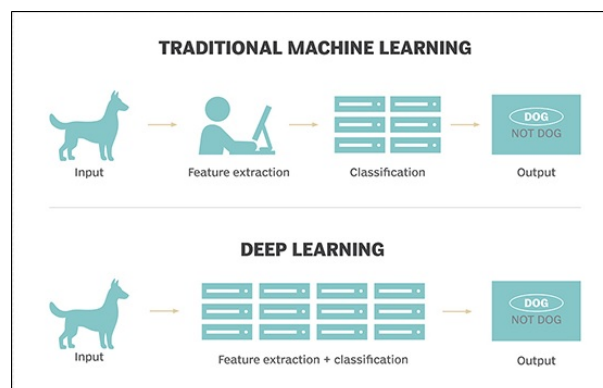


Figure 5.3: Deep learning

Source: TechTarget

This is a big advantage over traditional supervised machine learning because it does this automatically and it saves data scientists the time they would spend extracting features (e.g. ears, nose) or selecting the variables (e.g., long ears, short ears). Today, deep learning aids in image classification, language translation, speech recognition and it can be used to solve any pattern recognition problem with minimal human intervention compared to supervised training methods.

To use a DNN, a set of features are mapped out and assigned a random numerical "weight" which determines how the features respond to new data (e.g. images or sound). The DNN is provided with the correct answers. If the DNN does not accurately identify the input such as images or sound, then the system adjusts the weights (e.g. the importance each feature gives to the data given). After sufficient training, the neural network will start to recognize the correct patterns consistently.

5.4 Implementation of Machine Learning

One cannot just apply a machine algorithm and expect it to work. There are methods for applying machine learning for it to learn in the best way possible. According to Maarten van Someren, a professor of artificial intelligence and machine learning, it is considered an art when applying the appropriate machine learning method to a problem (Van Someren, 2001). Although some people might be more creative than others, there are proposed methodologies that can help guide this process. The following is a general method (Weiss & Indurkya, 1998):

- **Understand the problem:** To apply machine learning and choose the model class that is most appropriate, pick a model class based on the result sought after. This will help with identifying the right data.
- **Choose the right data set:** Finding the right data is critical. Data can come from one source or multiple sources. With ML, the more data we have, the better and the more ML has to learn. This data should fit the purpose sought after based on the model class chosen.
- **Clean the data:** This data must then be cleaned and optimized, which in turn creates a high probability of having an optimized machine learning algorithm. ML algorithms learn by digesting input data and outputs of that data that's best classified structured input and output pairs. Data can be clean by removing examples or values that don't fit the scenario.

- **Split into training and test set:** With a clean, structured and optimized dataset, the dataset can divide into a training set and a test set. The training set applied to a chosen algorithm to produce a training model.
- **Apply algorithm:** In this step, a training dataset runs through an algorithm to create a training model. The choice of the algorithm in this step is critical to get the right results. The training model must be scored and evaluated to see if it is good enough or not.
- **Score and evaluate the data:** Scoring and evaluating the training model helps indicate the performance of the trained model. For example, in the case of classification models, metrics such as classification accuracy, confusion matrix or logarithmic loss can be used.

The process might vary slightly depending on the learning style but holds true for the majority of ML implementation cases.

5.5 Considerations for using ML

Particularly in machine learning, understanding how these systems have reached their decision is very challenging. In the case of deep neural networks, hundreds of millions of connections all contribute to making a decision which make it very difficult to come up with a clear explanation. Machines are not capable of providing rational explanations for why one choice is selected over another, for example, either delivered an ad to a user or a specific medicine. Machines knowing more than it can tell is a case of reverse Polonyis Paradox. This is also called the deep learning black box problem. (Brouillette, 2017). In healthcare, for example, this seems to be slowing down the US Food and Drug Administration from approving new drugs and software that propose treatment or prevention. Drugs like Aspirin in the past were widely used for medicine but not well understood for 70 years. Lithium is another example whose exact biochemical mechanism affecting mood is still known as a drug approved for treatment of bipolar disorder (Knight, 2017).

The distinction between “out-of-sample” and “out-of-context” is an important concept to learn in ML and also a challenging one (Yeomans, 2015). Out-of-sample accuracy means new data generated from a static environment in an ML model will be able to predict the desired outcome. A clear understanding of the environment assists in defining and identifying the “signal” (a consistent relationship that we want to learn) from the “noise” (random correlations that won’t occur again in the future). All data has

a mix of signal and noise, and well-structured environments will help to sort through that to make better predictions. Algorithms draw their power from comparing new data to similar data from the past. When a model is applied in a different context, the cases in the database may not be similar, and now the strength of the model is liable to worse predictions. An example of this is an online store using data from online purchases to build an ML model for predicting new customers. If the same model was used to predict sales at a physical store even for the identical product, the sample model might not be helpful because it is out-of-context.

Even though some parts of the model building process can seem straightforward, it takes a vast deal of human judgment to figure out where the model will be most useful. There is much critical thinking that must go into a building in regularization and cross-validation framework. Many machines have biases that derived from the data provided to train models. For example, if a ML model uses the decisions made by human recruiters in the past as a source of truth, it may inadvertently learn their racial, gender, ethnic or other bias. Along with hidden bias, many machines deal with statistical truths rather than literal truths. Data that is not represented properly in ML models make it difficult to prove it will work as expected. For example, this can be concerning in high-risk or critical situations such as in hospitals.

Errors, in general, will occur, and one drawback with ML systems is that it is challenging to understand what went wrong. As seen, many of the solutions can get very complex to comprehend. Managers who plan to implement ML for decision-making should understand these risks and that in the ML world perfection is not the benchmark but the best available alternative. The advantage of ML is its ability to improve over time and give a consistent answer when presented with the same data and with the right person and technical skills. ML can be a handy tool for decision-makers trying to make sense of the inherent problem - hopefully without creating a new issue along the way (Yeomans, 2015).

DL provides significant advantages over earlier generations of ML algorithms by making use of larger and better datasets. However, with larger datasets, this means more processing power which requires supercomputers or specialized computer architectures to run and lots of time to learn. Processing power and time can become relatively expensive for large enterprises looking to make an impact.

6 Impact and Success Factors of AI and ML

The full impact of AI is still unknown, but its recent performance across all the great industries indicate the impact will be extremely disruptive. AI technology will create transformational impact such as all general-purpose technologies had done before. The steam engine, electricity, and the combustion engine all were catalysts to innovation and new opportunities. The slow or fast beneficial and risk impact will lay on the shoulders of management, their ability to implement and their business imagination.

To accommodate the disruption and transformation, we as people will need to be flexible and adaptable to the changes. For companies, to stay competitive and grow, they will need to learn quicker than before to survive. They will need to prepare for this by learning the new and discarding the old as challenging as that may be. They will need to change entire core processes, redefine jobs, and restructure the way productivity is measured. Not preparing for these changes can cause what Joi Ito refers to as "Whiplash." Joi Ito says, "If society can survive the initial whiplash when we trade running shoes for a supersonic jet, we may yet find that the view from the jet is just what we have been looking for" (Ito & Howe, 2016).

How is one supposed to drive on the fast paced AI highway? Many don't know whether they want to be in the driver seat in pole position (a first mover) or stay closely behind (fast follower) when facing the transformational development (Anthony, 2012). There have been positive examples of both with Apple not being the first mover in the digital music, smartphone, or tablet computing categories and now they are. Both models have been successful. Amazon was the first mover with electronic books and cloud computing, which turned out to be a success story. On the other hand, Myspace was the first mover in the social networking but then saw their demise to the simpler and easier platform Facebook . Both situations are plausible, but one important thing to remember as Scott D. Anthony puts it is that "ultimately, no one remembers who leads a race at the halfway point. They care about who crosses the finish line." He proposes not to ask "Should I go first?" but instead, ask "How do I accelerate the path to a

breakthrough idea?” (Anthony, 2012). In the case of AI technologies and applications, we need to ask which idea or technology will be the breakthrough idea.

To answer this questions, Alan Kay, the head of research at Xerox’s proposes that “The best way to predict the future is to invent it” (Anthony, 2012). He does not see much value in having endless sessions with experts providing prognostications; he sees more value in getting out and trying something. This approach helps with answering two essential questions such as what things have to be true for a transformational trend to play out? What experiments can you run to separate fact from the flawed assumption? Through these methods and mindset will leaders be able to adapt their business models and take competitive advantage?

The status quo of dividing up work between humans and machines is falling apart quickly (Brynjolfsson & McAfee, 2017). Those companies who stick with that mindset are going to have a significant competitive disadvantage compared to those who put machine learning in the places where it can integrate effectively with human capabilities. Access to these technologies will not separate the winners from losers in the business world; it will be those who are savvy enough to put this approach into place. AI will not replace managers, a manager who uses AI will replace those who do not.

6.1 Impact of AI on different industries

In general, the most frequent observations on how machine learning is impacting across industries have been in the following areas: activities and jobs, business processes and business models. In the case of activities, we see examples of machine learning based image analysis is helping identify potential cancer cells (Esteva, 2017). In the case of business processes, we see workflows being reinvented and redesigned with Robotic Process Automation (RPA) and optimization algorithms to help logistic companies with fulfillment like Amazon . With business models, we can see companies like Netflix who changed their business model from offering single DVD movies to a subscription service that using machine learning to personalize and recommend movies .

As we will find out in this section, one recurring theme is that machine learning systems are not replacing entire jobs, process, and business models and that the human factor is still highly relevant. Machine learning systems are augmenting human activities and making their work more valuable. For example, many chatbots today are not built to take over the support conversations altogether, rather they are helping advise salespeople improve their performance by recommending answers that have worked

in the past. McKinsey research shows 40 percent of the time traditionally spent on sales activities can now be automated using intelligent sales tools, making sales staff more efficient(Chui, 2015). The ability to remove repetitive tasks through chatbots help sales staff to focus on doing their job such as answer the complex questions and understanding the nuances of human emotion to close a potential deal.

ML impact is reaching many industries and areas of business. Those companies that have been able to hone in on its capabilities have been very successful in applying it. Here is a list of how real companies using it to expand growth, optimize processes and increase customer satisfaction:

- **Customer Service:** Chatbots are a significant contributor to the boom in adoption by companies who are looking to improve their customer service while lowering cost. Around 44 percent percent of the US consumers prefer chat bots (BI, 2017). Chatbots lead by machine learning algorithms that can process customer data, natural language and learn from each interaction over time are producing high-quality service without breaking a sweat. For example, the London National Health Service (NHS) testing an AI-powered chatbot on the 111 non-emergency helplines to 1.2 million residents. Callers can opt for a chatbot rather than talking to a person on the 111 hotlines. The chatbot, created by Babylon Health, will encourage people to share symptoms to consult an extensive medical database. The will can then receive tailored responses based on the information they have entered.
- **Marketing:** A New Zealand-based media company called Plexure used Microsoft's Azure Stream Analytics to analyze McDonald's big data (40 million+ endpoints) across Netherlands, Sweden, and Japan in the cloud to hone in on customer behavior patterns and responses to offer more targeted ads. Their work with the platform enabled them to target people on a sunny summer afternoon, walking near a store with an offer for a free ice-cream with a sandwich purchase. McDonald's in the Netherlands saw a 700 percent increase in offer redemptions, and customers using the app are returning to stores twice as often and spending 47 percent more .
- **Insurance companies:** Machine learning is helping to review with great speed and depth, thousands of insurance policies. It is providing the ability for insurance companies to review policies and calculate insurance policy payouts. Japanese insurance firm Fukoku Mutual Life Insurance is using machine learning to do this, and in doing so, it replaces its 34 employees for IBM's Watson

Explorer AI (Gibbs, 2017). The company expects to increase productivity by 30 percent and save around 140 million yen (£977,000) a year in salaries.

- Healthcare:** McKinsey estimates that pharma and medicine industry could reach a value of up to \$100 billion annually with big data and machine learning. Health care would benefit by having better decision-making, optimize innovation, improved efficiency of research/clinical trials, and new tool creation for physicians, consumers, insurers, and regulators (Benedikt F. C., 2013). SkinVision, for example, is a “skin cancer risk app” that lets users submit images of dermatologists in exchange for a personalized treatment plan. ML has also been applied to help monitor and predict epidemic outbreaks such as malaria by taking into account data such as temperature, average rainfall and a total number of positive cases among other points of data (Sharma, 2016).

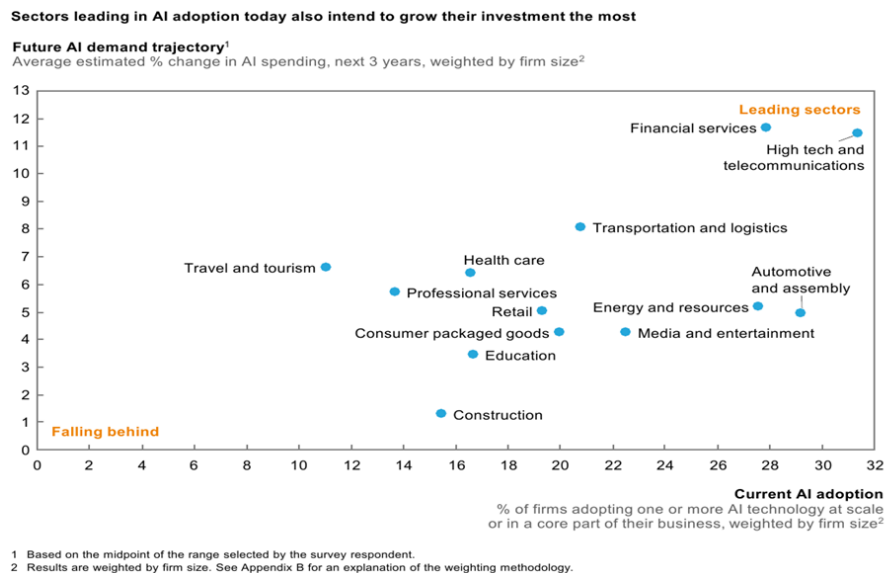


Figure 6.1: Sectors Adoption

Source: Mckinsey Global Institute AI Adoption and Survey (2017)

The ability for all the industries and sectors to grow on their ML adoption hinges on technical, commercial and regulatory challenges. As seen in figure the figure above, out of all the sectors surveyed by McKinsey, they expect finance, telecoms, and transportation and logistics to be on the edge of AI technologies in general. Many of these sectors have a high degree of digital adoption with high technical feasibility to put together compelling business cases. For example, the financial services and telecom industries have generated and stored large data sets which are structured, but this is

not the case for other industries like construction and travel.

When it comes to commercial benefits, those industries with complex business operations, where performance metrics are forecast based, reliant on fast and accurate decisions and personalized customer connections, are leading the adoption of AI technologies at scale. The finance sector will benefit from accuracy and speed of data. In retail the benefits can be identified throughout the entire value chain from inventory forecasts to automated customer operations and personalized marketing campaigns. Despite the clear commercial advantages, regulatory and social barriers raise costs and decrease the rate of adoption.

Liability of AI technology is a big concern in industries such as automotive and health-care. Automotive industries live from ensuring that their products meet all the standards to ensure the safety of the driver and all others in the surrounding environment. There are also privacy considerations that restrict data access, ethical issues that can create trained biases and job security concerns which are limiting market growth.

All these factors combined will determine the ability for AI integration to happen across industries. The first movers to leverage their technical skills, digital expertise and data resources are expected to be those with the most digital experience.

6.2 Dimensions of impact: Strategy, Organization and Technology

With the coming of the fourth industrial revolution, the changes to be expected will be significantly different than others before (e.g. Industrial Revolution). It will be non-linear, changing many paradigms and recreating them by gathering some together. This will require countries, companies, sectors etc. to consider many of the oncoming impacts holistically (Schwab, 2016). They will need to work together to make the necessary changes and to survive the hypercompetitive, changing and unpredictable environments that will come. AI will be one of those change drivers. Other examples of changing paradigms are developments in advanced robotics, 3D printing, internet of things (IoT) and sharing economy. To prepare for these changes, new organizational concepts and economical systems will need to be understood and renewed accordingly. The adoption of AI and ML will require machines to be used more and more in business, be it either robotic or computer based and it will be important to not “race against the machine” but “race with the machine”, as team mates complementing each other

(Zincir, 2016). To accommodate this, skills and managerial styles will require agility, sustainability, collaboration, flexibility and a sense of innovation through creativity.

In the following sections, this paper talks about the impact AI and ML will have on managing a business strategy, redefining people's roles in the organization and adopting technology while considering cost implications. In addition, each impact will be supported by best practices and supporting advice to support decision makers to successfully prepare the integration of AI. With decision makers in focus, these three dimensions were chosen because they can be highly influenced by top management and are interconnected.

6.2.1 Strategic impact

Strategy is defined as a plan to maintain and build competitive advantage (McGrath, 2013). On the contrary, it's been argued that "execution" not strategy offers an exclusive competitive advantage (Lippit, 2007). Now however, in this new fast paced landscape, competitive advantage often evaporates in a year, and if companies want to stay competitive they need to constantly start new strategic initiatives, exploiting transient advantages at once such as those from AI technologies . Transient advantage is defined as a strategy that prioritizes fast innovation to build a pipeline of competitive advantages (McGrath, 2013). Those who follow the transient advantage strategy can cope with emergent strategies easier than others. This is needed to stay competitive against customers that are very unpredictable and where fast paced competitive contexts are a normal . The decision makers leading the AI transformation can take on this strategic mindset and use it to help start filling the pipeline with these new strategic advantages.

To do this, they will need to understand two primary things: firstly how AI can transform the current business at hand and secondly, how they can walk before running with it. It will be their responsibility to create a clear vision, objectives and KPIs to measure the success. AI has the potential to create value and it will help companies win, serve and retain customers while staying competitive and growing. It is critical that decision makers develop a strategic plan that have the foundational elements in place for integration and will need to be ready to contend with all the ethical, political and social issues that will be also associated with it.

Back when big data became its own hype bubble, companies had to understand the impact of integrating data infrastructures that could collect all the data into one place

and query it to do analytics and answer questions about the business or product. As time has progressed, new software made collecting data more accessible than ever before. Organizations needed a data strategy to understand how to fit that into its agenda.

Now that data is more accessible than ever before, companies find themselves trying to make the best use of that data possible. As we know, ML is a perfect way for discovering, predicting, detecting and classifying data and using that insight to infer a plan of action. There are three common areas in which machine learning can impact the strategic decision making of the business and exploit existing data capabilities of a company:

1. **Understanding the business value:** The role of machine learning is increasing in influence in especially enterprise software applications that provide in-depth data analysis and adaptability such as in analytics and business intelligence tools. Business intelligence tools promises have pretty much stayed the same since its introduction by H.P Luhn in 1958, which is to collect an organization's data from disparate sources and process it in the best possible way to produce information that helps decision makers to make more informed decisions for the organization. BI systems do provide not only high-level decision support at a strategic level but also to middle management and operations. There has been an increasing need for BI systems to evolve and deal with the increasingly complex data analysis problems. Applications within BI systems need to not only deal with larger sets of more complex amounts of data and provide current status results, but they also have to predict and play with hypotheticals and make suggestions - the perfect opportunity for machine learning to shine.

Machine learning is the tool of choice for transforming the business intelligence applications approach into a wider intelligence or analytics platform. ML allows to extend the traditional questions BI answers like "what is going on with my business?" by giving a answers to questions like "why are we doing what we are doing?" and "how can we do it better?" or even "what should we do?". Machine learning for analytics and BI systems are increasing in popularity in areas such as risk analysis, marketing analytics, and advanced analytics. For example, a risk analysis team can use these tools to pick up fraud risks before transactions happen. Marketing teams can use predictive analytics to gain insight into the sales funnel from various target audiences, identify cross selling opportunities and reasons for a lead to fail to become a sale. This way of exploiting analytics

is on an entirely different level than achieved by other BI tools such as SQL reporting and Tableau. With ML, it is possible to construct value through many different sources of data never tapped into before - be it pictures, speech, text or other unstructured data sources. This surplus of data inherently helps ML to understand humans.

A BI tool currently bringing value to companies is SAP's cloud platform Hana, which manages databases by ingesting structured data, such as sales transactions or customer information from all various access points (mobile, desktop, sensors, product equipment). For example, if sales teams use company smartphones or tablets to record purchase orders, data can be used for those transactions to be analyzed, understood to spot trends and irregularities. Hana can be accessed through servers or via the cloud. Walmart uses Hana to process transaction records (across 11,000 stores) within seconds to speed up operations and control back office costs by consolidating the processes and resources needed to handle the work (Daniel, 2016). Hana is also being used to call attention to various parts of the machinery that are running slower. For example, a factory manager can monitor the equipment on an assembly line and using Hana, the results from the products can be queried to determine if a new course of action is needed such as an inspection or a replacement of a piece.

SAP is not the only one providing businesses with machine learning platforms for business intelligence. There are also companies like Domo, which uses a cloud-based dashboard to gather information from Salesforce, Square, Facebook, Shopify and other applications to gain insights into customers, sales or product inventory. Companies like MasterCard, Univision, eBay and SAB Miller are planning to use the platform. Univision, a television broadcaster said they use Domo to connect to Google Analytics, Facebook and Adobe Analytics to get more value from programmatic advertising .

2. **Product development:** ML has the potential to aid in product development and engineering planning. These departments now have access to use machine learning to improve products in a more real-time than reactive way. ML for product development can help to improve current products or innovate new products. These products do not need to be necessarily client facing, but it can also be internal software the company can use.

Here is a comparison of new ideas that were successful in the past and ones that are being reinvented today with ML:

Before ML	After ML
Customer relationships	Automate follow ups and replies for customer relationships
Simple UI for managing photos into collections and albums	Automatically organizes photos and the ability to search by face, location, thing...
Crowd sourced recipe database	Recommends what to cook based on interested revealed in social media feeds
Easy to use website and app for tracking my finance and budgets	Analyze my expenses and tell me how my budget compares to other people with similar income, tell me most effective way to save more money

There are some best practices in utilizing ML for product development purposes. The process starts by understanding what needs to be solved, what is the business problem and what is the goal of the product. In the ML world, in most cases one algorithm feeds into the next one creating complex relations where for example an algorithm is figuring out the value of a given user and that value is being used as input for another algorithm to determine what to recommend. For example, products like Quora or Netflix use dozens of different algorithms, and there will be many ways to combine them. One way to not over-complicate things too quickly is to start with a simple goal until it has been reached and then build on top of that based on the feedback and functionality it provides to the end user. Introducing functionalities based on this principle will help to assess where new ML is working.

Xavier Amatriain, VP of engineering at Quora and previously a research/ engineer at Netflix, recommends that when ML is considered for product design, following three things should be done (Amatriain, 2016):

- Work on the ML algorithm that will get the desired answer based on the problem
- Work on the product design that is going to make sure of that algorithm
- Work on a system that will support and scale that algorithm and the large amount of data with it

The key to success in the product creation is integration between machine learning algorithm, the product and the system behind it. For this to work out, product managers and leaders should set the right expectations. Many product leaders want ML not just to analyze data but to think for them, for example, tell them why a customer is not converting or not doing what it is supposed to do. However, that is not the job of ML. ML is there to help with making predictions and getting to the goal more efficiently with greater accuracy, but not to explain why. Predict, a product that provides insights into which users will convert or not before they do is an example of an ML supported product built by Mixpanel, a mobile analytics platform. The product lets companies select an action (clicking a link or downloading brochure), and based on a user's previous behavior, the platform will grade the users on how likely they are to complete it. Then the brand or company can target their users based on the best grade (Mixpanel, 2015).

Another perspective on building AI into the product strategy is not so much the design or creation but the maintenance of products which is seen more and more in the manufacturing industry. Most manufacturers are using ML to inform engineers with maintenance, repair, overhaul (MRO) and performance predictions to a component level which helps to save costs. For example, Mueller Industries used Augury who specialized in predictive maintenance to extend the operational life of machine components, eliminate damaging equipment failure and reduce costs of parts and labor. Using vibrations of the mechanical systems, they were able to use their diagnostic technology to analyze machines with a portable device called the "Auguscope", upload the reading into the cloud and use Augury's machine learning algorithms to analyze the data against a library of signals to then immediately provide feedback to the technician.

There are many product development examples with ML, and companies are still figuring out how to benefit from them if they have not started to do so yet. Be it an internal product or an external, the advantages of ML can help leaders to make decisions that can ultimately contribute to steering their short-term and

long term strategies. Innovation, however, does not happen overnight which is why R&D plays a major role in experimental ML.

- 3. Research and Development (R&D):** Not all companies have an R&D department in place, but in some industries it is the heart of a company and focal point for innovation. ML helps to improve R&D capabilities using data to come up with new products, new business ideas, new revenue opportunities, etc. Global competition, high costs of product development, changes in compliance and the need to go fast to market places lots of pressure on R&D departments. In industries like aircraft and automobile industries, the R&D is the epicenter of an organization. They are looked at by the company for developing cutting edge products and services that keep the company growing and relevant. For R&D to be at the forefront of innovation, they need visibility into data across the entire business from customer, expert, scientific product information to market, regulatory and competitive information. The more data, the more opportunity and challenges, which is why machine learning is a useful solution for R&D departments to harness insightful information and drive innovation on potential products or company solutions (Rndmag, 2017). R&D departments, however, are not always understood fully by their leadership team. R&D teams needs leadership that understands experiments that don't work are just as important as experiments that do work. With this knowledge, companies can avoid technologies that haven't proven their importance and cause even more damage down the line if they are put either into production or released to the open market. In general, it is important that companies pursue understanding business value and product development before humping into R&D. The practice of investing in cost-saving projects is easier to sell internally than new revenue opportunities in certain cases. There needs to be a balance between R&D and the rest of the organization's goals to stay ahead.

It is the responsibility of top management to prepare their organizations for AI by finding the best place to start. In most cases, they will need to decide whether they want to build AI applications inside the core business proposition or enhance existing workflows. Wherever they plan to start, identifying and understanding the problem they want to solve or client expectations they are trying meet will be essential. In most cases, articulating this issue is a challenge. For the enhancement of workflows, it should be narrowed down to either reducing the cost or increasing revenue based on being able to get the right data. For example, if a problem reduces churn rates, the right data could help detect users who are at high risk of leaving by analyzing their

activities on a website, SaaS application or social media (Altexsoft, 2016). While there might be metrics that already can provide red flags, ML algorithms can assist with discovering hidden dependencies.

In many cases, ML works best where standard business logic and rules are not sufficient to solve it. ML can help where decisions rely on subjective opinions by either the expert or the decision maker. ML is very helpful for people to see through their biases based on real evidence. Although ML can support, it should not be a mean on its own. Management will need to truly understand the limitations (positive or negative) to leverage this technology the best way possible. A vital role that has driven the successful education and transformation of technologies in many organizations is that of the CIO.

Evolution of the CIO and rise of a CAIO

With the evolution of IT, the internet and boom in data, the CIO had a central role to the company's strategy organizing central business tools, laying the technical groundwork and connecting the pipes to all sources of possible information. Historically, their main responsibility has been to collect business requirements and then select and deliver the right application for it. Until around 2005, the biggest transition CIO's were responsible for were going from mainframes to client-server technology until the introduction of cloud based service applications for enterprises, small data to big data, unintelligent applications to AI powered applications democratized the power and altered the control.

The cloud enabled the ability for data sharing across company silos. The impact of the cloud was critical for the CIO for two major reasons. First, it switched much of the business capital cost to a variable operating cost where the premium was on speed and agility. The second reason was that it changed the balance of power, from there being six to eight major business applications for the entire company to have access to over 4,000 applications available. This required the CIO to develop new policies and frameworks. The abundance of tools and applications available was the catalyst to the generation of much more data.

As occupied CIO's managed and enabled applications that were in line with the business process, the concept of data became very valuable and also demanding to manage. The developments in web technology, social media mobile, and sensing devices were a huge catalyst to the explosion of data. For example, Twitter processes over 70M tweets

per day, generating 8TB of daily data (Groginger, 2017). This introduced CIO's to the V's of big data: variety, variability, velocity, and veracity. The importance of storage and computing approaches had to be understood, and analytic capabilities had to be built to present insights on the data collected.

Despite the troves of data and cloud based applications available, they are useless unless someone could take action on it like update a sales opportunity or program a defined rule to follow. Many of these requests were backward looking, using historical data, only using what already happened to react to future events. Looking into the past, data was not enough, and there was more pressure to gain more value it. This lead to applications becoming intelligent and the increasing demand push for forward looking insights like predictive analysis. The first companies to explore intelligent applications were in the marketing and sales, sectors such as Quantcast with predicting audience targeting and Everstring with predictive lead scoring.

This shift in demand had a significant impact on technology players, so they started to integrate the intelligent use of their data as an integral part of their core business proposition. For example, Google's CEO stated they are now an "AI first" company, and Microsoft has articulated its "ten rules for AI". Many S&P 500 companies, however, are still struggling to fully adopt a digital strategy.

AI is still immature and evolving quickly, so it is unreasonable to expect everyone in the C-suite to understand it completely. If an industry generates large amounts of data, it has a high probability that AI can be used to transform that data into value (Ng., 2016). It is only a matter of time before AI hits the enterprise large scale accessibility as it did with cloud based applications and data before.

For those companies that want to get ahead of the curve and have large sets of data but lack of AI knowledge, it is time to create a new top management role, such as a Chief AI officer (CAIO) or a VP of AI. Only by putting decision makers in the C-suite of organizations into place who would be responsible for overseeing the integration of AI, it can become an integral part of a company's strategy. Building a comprehensive strategy on AI will need their understanding and support to communicate it to the organization. A leading expert, Andrew Ng, states that to succeed in the future, businesses need to appoint a CAIO. Those who land in this position will need thorough training in applied ethics that prioritizes the values of users and computers (Ng., 2016).

The challenge for companies will be to find a qualified CAIO and talented team to support. AI specialists are scarce and depending on the company and job, it is challenging to attract top talent. According to McKinsey, by 2017 the demand for expertise for data scientist will be 60 percent higher than supply. One of the reasons for this scarcity of qualified data scientist is the required skill set that combines math, statistics, programming, databases and domain expertise. A CAIO and a dedicated AI team are critical and provide a higher chance of successfully leveraging AI benefits for the company. An AI leader can dramatically increase the odds of success, depending on the person. Although data scientists should be ready to educate decision makers on opportunities and limitations of ML, most have a hard time translating complex results into concrete business language. The CAIO should possess a special skill set that allows to understand the data scientist, to see the big picture, to see the options to scale ML application and combine the business and technical vision. To help narrow this skill set down here is a list of traits recommended for a potential CAIO or VP for AI (Ng, 2016):

- Technical understanding of AI and data infrastructures
 - Should have built and shipped nontrivial machine learning systems
 - Understand how to organize databases and make sure they are stored securely and accessible.
- Ability to work-cross functionally
 - AI is not a product or a business. It's a foundational technology that can help existing lines of business and create new products or lines of business.
 - The ability to understand and work with diverse business units or functional teams.
- Strong intrapreneurial skills
 - AI creates opportunities to build new products, from self-driving cars to speakers you can talk to. A leader who can manage intrapreneurial initiatives will increase your odds of successfully creating such innovations for your industry.
- Ability to attract and retain AI talent
 - Needs to know how to retain talent by emphasizing interesting projects and offering team members the chance to build their skill set.

In case a leader for AI with this outlined skillset can't be found, people within the organization need to be identified who have technical know-how and strategic foresight

and work directly with problems. Machine learning frameworks, as we have seen with Microsoft Azure, can ease the integration process provided that the data is available and goal of the task at hand is defined. In many cases, finding a solution with traditional ML is possible - as long as the data at hand is structured, well understood and labeled. For example, using a decade of transactional data can use machine learning to find correlations between customer's demographics and products (Hammond, K. 2017). In conclusion, the key is to have a team in place which can communicate business problems and understand the technical solution and data. Having a leader to represent this team will definitely help to drive success.

Preparing for strategic AI integration

Many companies who seek to integrate AI tend to experiment with AI without a holistic strategy in mind. These companies do not have CAIOs yet, and neither do they have a strategy. The pressure to do something is there. However, integrating AI randomly can be costly as we will see more in the technological impact section.

To make sure the top management can execute their strategy properly and balance the opportunity of AI technology while mitigating its risk, a list of key factors to build an AI strategy plan is described. This list is based 612 global respondents who participated in the the Forrester 2016 Q 2: Global State of Artificial Intelligence survey which can be used by all companies and industries who are working on artificial intelligence and automated systems (Gaultier, 2016):

- Hire an AI Lead
 - Either it is the CIO, CEO (depending on technical background) or CAIO - there should be someone in the position to lead all AI related activities. Good AI will depend on good data stewardship. The person in charge should understand the tradeoffs from both a technology and non-technology perspective. It is important to have a clear plan on why and how AI will be introduced and integrated. Dependent on the plan, the organizational impact and change management impacts will vary.
- Build a roadmap to prepare for wide-scale AI deployment
 - One of the first steps the lead decisions makers should take is to assess their organization's current state and if it is ready for AI. For example, are the company's data assets in good shape? How much do current decision makers leverage data and analytics? If the organization is not currently

driven by insights or in the process of doing so, an integration of AI is likely a waste of resources.

- Develop AI talent strategy
 - Since AI will be new to the organization, all the skills needed will not be readily available, and with a shortage in general in the market, it will be costly and hard to rely on external resources to fill in those gaps. A decision maker will need to rely on all options such as:
 - * How can we train the people we currently have in the roles we need?
 - * If it makes sense to hire external resources, would a long-term partnership make sense or would a more educational partnership?
- Take an end-to-end process view and extend across processes
 - AI will change the way that processes are structured and redefine jobs. These processes and the profiles involved will need to be reviewed. A process alone can be automated, but the true benefits will be seen when automation is done across processes.
- Set up governance model to mitigate risks
 - All known issues must be up for discussion, and there must be an agreement in place on how to deal with all the known possible implications AI might have. Either it is privacy, liability, discrimination, there must be rules set up and governing principles for AI systems.

The introduction of AI will imply significant transformational change on everything from customer interactions to back-office operations, supply chain and manufacturing. It will be important to understand that machines are not going to start thinking on their own with evil intentions to unleash destruction. Humans will need to understand how to judge best the use of AI and to what degree they can trust it to make the appropriate decisions where it is introduced. To build confidence is a very critical human element that could make or break the deployment of an AI strategy.

In conclusion, it is important to demonstrate strategic flexibility, the ability to precipitate intentional changes and adapt to environmental changes through the continuous rethinking of current strategies, asset deployment and investment strategies. The key ingredient for strategic success and creating value that can create real impact is taking a “people first” approach. To make this transition with AI, companies and governments need to acknowledge the challenges and change how they behave. People will

need to be prepared - intellectually, technologically, politically, ethically and socially. This could either be lead by a team or the appointment of a CAIO. After knowing the why and how of AI integration, the organizational impact - the impact on the people - will be in focus.

6.2.2 Organizational impact

The impact of AI on the role of people, the work to be done and how people are measured will transform the foundational concepts of an organization. From the recruiting, the nature of tasks to the wages, company culture and organizational hierarchy, the integration of AI is set to challenge all of these elements. It will be the responsibility of management, HR and the employees themselves to redefine jobs and processes affected so that the organization as a whole can take advantage of AI's automation potential. The ability to properly staff, manage and lead these organizations would bring benefits in the form of increased output at high quality and reliability as well as performance speed which could generate financial benefits.

Contrary to popular belief, low skill, low-wage roles are not the only jobs which will be susceptible to automation. High paid occupations such as financial advisers, doctors, and top executive activities will also be exposed to automation (Chui, 2015). For activities to be automated, processes must change and the jobs must be redefined. These useful AI technologies will not integrate themselves and they will not "think" on their own either. The human factor will still be very critical for the near to mid-term of these technologies, and we need to make sure the human component is ready to work with its machine counterpart, assistant or extension to bring out the best in both.

Automating activities

The automation of tasks is upon us, and the success metric will be based on how well we learn to work with machines to augment our capabilities and increase productivity. Automation will not be straightforward, it will depend on the complexity of the activities involved in doing the job. However, the takeover from a workforce of machines and algorithms is a real concern, and as we see with disruptive technologies such as Uber (displacing taxi drivers), Netflix (disrupting the rent a movie business) or Airbnb (disrupting the hotel industry), it is a concern to be considered and actively addressed in the workforce.

As learned in the chapters before, routine tasks are top candidates for automation via AI. However, many technologies can already match or even exceed human performance required in more complex tasks (Chui, 2015). Companies like Narrative Science already analyze data to generate natural language, being able to write reports in seconds. Amazon’s Kiva robots with automation technology can plan, navigate and coordinate warehouse orders and IBM Watson can suggest available treatments for specific ailments, based on medical research for specific diseases.

Currently, the way organizations can look at different automation approaches and opportunities to integrate AI are (Davenport, 2016):

- **Robotic process automation (RPA):** RPA is for high volume, low complexity, and routine administrative “white collar” tasks. RPA supports in is a successor of outsourcing many administrative processes, reducing costs and increasing accuracy. Tasks such as requesting customer identification information and tracking the status of a delivery are optimally done with RPA.
- **Cognitive automation:** Cognitive automation takes on complex tasks like pattern recognition or language understanding. For example, the “Amazon Go” retail store has no cashiers or checkout lanes. Customers pick up their items and leave the store, sensors and algorithms automatically charge their Amazon account. In this case, automation has replaced work elements of scanning purchases and processing payment. Other elements of the “job” of the store associate are still done by people, such as advising in-store customer about product features.
- **Social robotics:** Social robotics involves robots moving autonomously through sensors, AI, and mechanical robots. For example, “driverless” cars, where robotics and algorithms interact with other human drivers to navigate through traffic. When deconstructing this, it reveals that the human still plays an important role. A human “co-pilot” no longer does the work of routine navigation and piloting but still does things like observe the driverless operation and stepping into assisting with unusual or dangerous situations (Boudreau, J., Raven, J, 2017). What is not explicitly said is that the human co-pilot is “training” the AI-driven social robotics every time the human corrects the situation which in turn helps the AI system to “learned” (Boudreau, J., Raven, J, 2017).

It is important to highlight again, that not all automated tasks will reduce labor costs or displace labor in general. Many of the tasks through AI application such as machine learning algorithms are for example being used to reduce factories energy usage

through predictive maintenance or to optimise logistical routes for more efficient distribution. Also, the McKinsey survey has shown that 24 percent of firms expect the size of their workforce to increase in response to AI due to the anticipation of job growth by more efficient activities . The nature of the jobs can change.

Automation adoption will vary per industry and only if the value proposition makes sense. A job with potential to be automated doesn't necessarily mean that it will be. Relative costs are a key component in evaluating an automation option. Industries will benefit from the adoption of automation through increased process efficiencies, labor cost savings, and innovative products. Although barriers have been lower than before with open-source platforms, industries will still face high costs of research and development associated with the next technology. For example, Toyota Motor Corporation recently announced an additional 50 million dollar investment in robotics and A.I research supporting its efforts to develop intelligent cars. (NYT, 2015)

Redefining jobs and engaging resistance

It is said that low-skill, low wage activities will be the first ones to be affected by automation. However, research shows significant percentage of activities can be performed by even those in the highest-paid jobs as well. An estimated 20 percent of CEO time could be automated, such as reviewing and analyzing reports, data to make decisions and drafting assignments (Chui, 2015). Regardless of low or high wage, there are still many jobs with activities that cannot be automated with technology available today. According to a research analysis, fewer than five percent of occupations can be entirely automated based on today's automation technologies. However, 60 percent of occupations could have 30 percent or more automated.

This significant impact will require redefining some jobs and transforming business processes to account for the automation.

Redefining and automating will require the ability to deconstruct work into components. Deconstruction then reconfiguring the elements within the job will help reveal the human-automation combinations that are more efficient, effective and impactful.

For example, mortgage-loan officers will spend more time advising clients and setting expectations then reviewing and processing paperwork. In health, doctors will be able to focus on more rare cases and improve accuracy of most common issues by saving time making the diagnosis. Lawyers can comb through thousands of documents

through text-mining techniques and identify the most critical documents that need to be reviewed, saving time and extra needed resources. Sales organizations can increase time of their salespeople interacting with customers by automating the ability to generate leads and identifying opportunities for cross-selling and upselling(Chui, 2015). These are common examples of machines augmenting the human capability and capacity to focus on things of higher value. This will challenge organizations to retain human work. The reconfiguration of non-routine activities will yield new and different types of jobs. The jobs will require broader and more creative thinking, things human exceed compared to AI today.

Creativity and sensing emotions are difficult to automate. These two things are core to the human experience but the amount of time spent on activities that require these capabilities is meager. Research shows that four percent of work across the US requires creativity to a medium level of human performance and 29 percent of work activities require a median human level of sensing emotion (Chui, 2015). The lack of creativity suggests potential to generate a greater amount of meaningful work where employees can focus more on utilizing creativity and emotion while avoiding routine and repetitive tasks. For example, financial advisors can spend more time understanding clients' needs and explaining creative options rather than analyzing their financial situation, or interior designers can spend less time taking measurements, creating illustrations and ordering materials and more time developing innovative design concepts based on the client's wants and needs.

Redefining roles will also mean dealing with resistance. Resistance to change is high on the challenges companies will face when implementing AI into their organizations. It will take time until companies build up cultural and intellectual assets where people know enough about machine learning and machine learning people know enough about the business aspects, and they can work together productively to get these applications to work. In large companies, structures are more rigid and not agile. It often takes years or decades to change each time a paradigm shift comes along as opposed to small companies like start-ups who can bring together people from different backgrounds and let them work together freely. This is not a new topic for top management or executives in decision-making positions.

According to an Accenture survey of 1,700 managers in 14 countries and interviews with 37 senior executives responsible for digital transformation, there were significant patterns in manager's attitudes towards AI that were identified. In this survey, they found that 84 percent of all managers expect AI to make their work more effective and

exciting, yet 36 percent expressed the fear of job loss (Shanks, 2017). It also found that the further away a manager is from the C-suite the less enthusiasm was expressed for AI. This means that top managers see the opportunity to integrate AI into the workplace, while mid-level and front line managers are less optimistic.

When asked about trusting advice from intelligent systems on business decisions, 46 percent of top managers strongly agreed with the statement, but only 24 percent of middle managers and 14 percent of frontline showed the same level of agreement (Shanks, 2017). A similar trend was seen on the question whether they would be comfortable with AI monitoring and revolution their work, 42 percent of top managers said they strongly agreed while only 15 percent of front line manager shared the same sentiment. This finding presents that executives cannot assume that mid and lower level manager share the same perspectives. These differences can cause trouble for all involved in adoption. Leaders who fail to account for this resistance with any AI strategy and try to implement it top-down will likely fail (Kolbjornsrud, V. 2017).

To make sure this does not happen, leaders will need to work on approaches that balance human judgment with machine generated advice. Each company is different, and there is no one-size-fits-all solution for this. It will be an organization wide effort establishing acceptance. For businesses that are global, with centralized methods of managing their operations, it should be noted that cultural differences in opinion should also be taken into account. The graph below illustrates that according to a survey, 46 percent of managers in emerging economies strongly agree that they trust intelligent machine advice for future decisions, but only 18 percent said the same in developed nations. This may be because emerging countries see this as an opportunity to achieve advantages against competitors (Bradley, et al. 2012) while the developed nations that are more mature in technology understand what it takes to make it work being more skeptical. Overall, this reinforces the message that executives cannot expect to use a one-size-fits all strategy of AI in organizations. Their adoption strategies should be customized to local and organizational nuances. Top management should make sure to think how the end user or employee, in this case, will use and benefit from the technology. For example, the introduction of intelligent schedule systems that match employee capacity should use algorithms that adapt to the employee's skills and learn their personal preferences to improve their scheduling effectiveness. With the employee in mind and getting them involved, there should be better adoption outcomes.

In an analysis on how organizations can win over skeptical managers, three steps to success are recommended:

- **Start exploring together:** Top executives should involve managers from different levels and geographies in initial experiments with AI technologies and in efforts to scale them within the organization. There is no set of rules or maps to help navigate together, which is why they need to explore together learning quickly and applying their feedback where possible.
- **Keep track of AI use:** As part of the integration plan, executions must define early on metrics that can be used to measure adoption. They can help top-management identify proper use versus potential flaws. This also involves taking special care of employee information that is sensitive such as data collected on their location and work habits that might be captured. Management should make sure to monitor this data and safeguard it from distributing it in ways that could hurt the trust or violate data-privacy laws.
- **Craft new recruitment and training strategies:** As these technologies are adapted, managers should make sure they recruit and train talented people who are willing and able to collaborate with such intelligent agents. This will put a premium on “soft” skills (e.g., collaboration, creativity) than technical competencies in the future (WSJ, 2016), which will complement the analytical qualities of the ML algorithms and intelligent systems.

The success of AI and ML in an organization will highly depend on how it is introduced. When done properly, it has all the potential to bring the organization forward through increased efficiencies, but it can also hurt the company culture. It will be up to top management to decide whether to sit and watch the impact of AI happen or actively shape it. They will need to decide the level of involvement they will want from the organization in adopting AI based on their judgment and goals. The overall consensus, however, that engaging with managers and the front line through exploration will help the organization embrace the promise of AI together.

The next section will deep dive into the impact on technology from a cost perspective. that management will need to take into consideration through their integration efforts.

6.2.3 Technological impact

Technology is a tool and in itself does not deliver competitiveness improvements. They must be planned for, understood and implemented correctly along with a team to support it. As we have seen, AI technologies play an underlying support role providing better-informed decisions which ultimately create business value in many forms for a company. Machine learning, in particular, has been growing in popularity due to its

continued ability to create value. Forrester describes machine learning platforms as an AI technology in its growth phase due to the growing number of implementations of ML that are producing evidence it is working. ML, although very powerful, is still a very complex system, which may require significant investment in time and resources to deploy depending on its use case.

On the top of the list of challenges when implementing AI technologies that impacts the business, is cost of technology. As seen below in a survey conducted by the economist, 27 percent of executives chose that the most practical challenge in implementing AI into their business was the cost of technology (Economist, 2016). Culture resistance was second with 22 percent, which was addressed in the previous section and finding a provider was 21 percent as for IT capabilities and data quality.

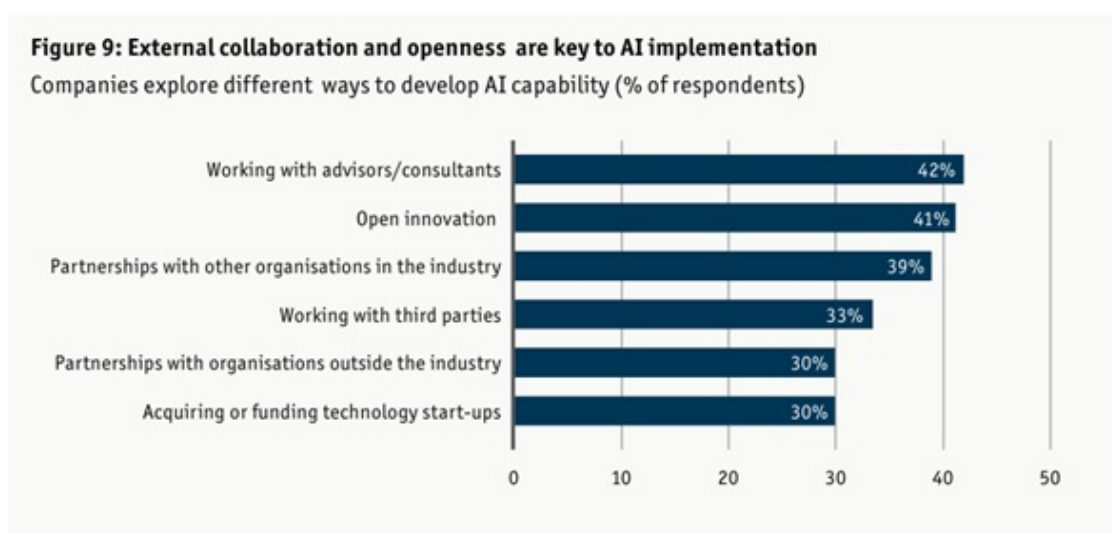


Figure 6.2: Companies Explore Different Ways to Develop AI Capabilities

Source: Artificial intelligence in the real world. Economist (2016)

When it comes to smaller companies, an AI solution can be easily managed and agreed on due to the size of the company (less bureaucracy and processes) by executing it through third-party cloud partners. Smaller companies also have the possibility to take advantage of open-source AI platforms from big technology companies. This allows smaller companies to train data sets much more quickly internally at affordable rates. Most of the software is also easy to use as well with a large community backing for support. For larger companies, using third-party partners can become costly. The higher hurdles on bureaucracy, security guidelines and number of involved people

causes longer integration periods and higher costs.

Cost of Data

Traditional ML algorithms and approaches are struggling with the influx of data. They were struggling with their volume, velocity, and variety. Big Data requires new processing paradigms to enable the insight discovery, improve decision making and process optimization. The ability to extract value from Big Data depends on data analytics. Data analytics involves various approaches, technologies, and tools such as those from text analytics, business intelligence, data visualization and statistical analysis (Jagdish et al. 2005). Machine learning is the primary driver of the Big Data revolution. ML can learn from data and provide data driven insights, decision, and predictions. Massive datasets impose challenges to traditional algorithms which were not designed to meet such requirements. Many ML algorithms were designed for smaller datasets, with the assumption that the entire data set could fit in memory. Another assumption was that entire data sets were available for processing at the time of training. Big data breaks these assumptions making traditional algorithms unusable or significantly impeding their performance.

When it comes to crunching these massive datasets for machine learning and powering predictive analytics, there are two types of major costs involved: the cost of training and the cost of predictions - which scale differently. AI is effective when it can analyze and learn from large volumes of data. Quality machine learning models will only be as good as the quality of the data, and anything less can be very costly. According to an IBM estimate, 3.1 trillion dollar is the yearly cost of poor quality data in the US alone (IBM, 2016). Reasons to why bad data costs so much is that decision makers, managers, data scientist and others must accommodate it in their routine work (Redman, 2016). This data handling is time-consuming and expensive due to the errors in data or the corrections of those errors people need to make to complete the task at hand. This error can then also leak through the system affecting the end customer either it be an employee who cannot find an order or a customer who did not get the right product delivered. These steps of data generation and correction loops are prominent in all types of organizations - companies or government agencies. These so called “hidden data factories” can be expensive. It is estimated that 50 percent of workers time is wasted on hunting for data, finding and correcting errors and searching for other sources they trust (Redman, 2016). Another 60 percent is an estimated fraction of the time that data scientist spend on cleaning and organizing data (Crowdfower, 2016). At the same time IT is expanding enormous efforts to line up systems that do not

communicate across silos which all sums up to non-value-added work for all of those involved.

Both, old infrastructure and bad data can be a challenge for some industries to adopt ML. For example, in healthcare many systems are fragmented and based on old models, and only a few have correctly embraced data collection. Research in healthcare through AI is difficult due to this challenge in accessing large medical datasets for legal or other reasons. The cost of prediction is the hardware and energy costs associated with a single computation and prediction. As AI technology deploys worldwide, the prediction costs will scale with every use. The cost of training is also high because you are dealing with people's lives and toleration for bad data.

In general, solutions for this are not as easy as cleaning data up. As with all other new best practices, it requires a new mindset. It is not always clear how data is being generated, where the errors are occurring and then which source is up-to-date with the latest data. The ability to improve data quality is an ongoing cost savings practice that will enable more opportunities to build on top of other data strategies.

Technical Debt

The need to build and ship new products and services can lead to trade-offs between speed of execution and quality of the engineering. This is a real world challenge faced often by real engineers. In 1992 the concept of technical debt was introduced by Ward Cunningham as a way to help quantify the cost of such decisions. The cost in most cases is paid for with refactoring, debugging and complicated testing.

Companies like Google, for example, have been running machine learning in production for a long time now and have identified some areas that accumulated technical debt (Sculley, 2015). Google explained their experience in a paper that outlined the debt problems and general best practices for implementing machine learning models without having to accumulate high level of technical debt. The paper concentrates on the fact that machine learning systems are fundamentally different from the development of traditional software and identified five areas that contribute to technical debt:

- **Erosion of boundaries:** Software design is typically modular. The modules isolate regions of related code that perform defined tasks and separate them from the other modules they interact with. This disentangling of code makes it possible to test code and for it to be maintained by different people. ML models, however, entangle input signals from various systems together which make it

difficult to see the boundaries for a specific intended behavior. The algorithms and data are so closely dependent, a change in the external data source would change the way we would expect the algorithm to behave. In most cases data acquisition, preprocessing and model tuning is likely to be managed by different people. This means it would be difficult to maintain changes to the underlying data sources being relied on. This is an inherent problem that must be accepted and prepared for.

- **Data dependencies:** ML have large data and complex feature dependencies. Any loss in one feature, input of a new feature or change in the values of a feature may alter the ML model entirely. This makes it difficult to avoid the CACE principle of Change Anything Change Everything. Data dependencies that are unstable or underutilized are difficult to detect in machine learning systems which can lead to developing forms of hidden technical debt (Sculley, 2015). Because dependencies such as features tend to be finer in ML systems, it is important to invest time in understanding all the dependencies and controlling them.
- **The feedback loop:** ML systems learn in real-time and as they learn their behavior changes. For example, search engines adapt their algorithms based on the click-data it receives from its users. The same effect is also apparent in machine learning-based recommendation systems. These feedback loops are similar to filter bubbles, in which people are only shown views and posts they are predicted to agree with which means it makes it difficult for them to interact with the different content they might not agree with controlling them (Lipton, 2015).
- **Dead experimental code paths:** Dead experimental code paths are code which was designed for experimentation but shipped into production. This could result in experiment code being selected that could mean disastrous consequences while not presenting any easily detectable software bugs.
- **Glue code:** Advanced ML algorithms are often implemented in packages that provide generic solutions. This often results in engineers writing glue code into these packages to, for example, process the companies data, set the values of parameters, select the appropriate algorithm, report results, etc. This presents a problem of vulnerability if any changes are made to the underlying external libraries.

For a data scientist, it is important to be aware of the complexity of the models they develop and impact these models will have on their organizations and how much it

will cost to maintain. ML models used in business applications interact directly with the real world, and since the real world is unstable, it can result in the accumulation of hidden debt within a ML system. This is why it is imperative for companies who choose to implement ML systems to create capabilities either internally or make it a prerequisite of their third party vendor that helps prevent non-obvious hidden debt created unintentionally.

Return on Investment

As we have learned before, one key challenge to get started with AI is having a reliable business case in place. The ROI is one of the key KPIs companies will need to consider when setting up AI.

Most gains can be categorized by either predictive insight or process automation (Archees, 2016). These two categories are either lowering cost or increasing revenue. In the case of predictive insight, companies can customize customer behavior which in turn provides more opportunity for sales. In the field of medicine, machine learning can help with anticipating medicine effectiveness to reduce time to market. In most user or subscription based organizations, forecasting can be used to see if they are about to churn to improve retention. Machine learning in this context helps increase reactivity by providing the tools and information to make decisions faster and with more accuracy. Process automation and efficiency see its ROI in, for example, augmenting the management of decisions, which in turn for people in the financial services sector can mean better margins and help mitigate costly mistakes. On the production lines, with increased precision, you can alleviate the need for more quality controls. In companies who rely on supply chain efficiencies, ML can help to save time according to user demand and costs during delivery.

On a macro level, according to a frontier research report, with conservative assumptions, the potential of AI applications is set to double the economic growth rate. This is due to the improving of human productivity by up to 20 percent. The economic impact of AI on specific countries are looked at through a metric called “national absorptive capacity” (Purdy & Daugherty, 2016). Which means, how quickly a country can absorb and spread the impact of innovation. The US and Japan for example score high on this due to their strong innovation in investment, education, learning, and ability to move technology quickly as opposed to other countries who have labor constraints or capital to inject into this type of projects.

Companies in countries with larger populations have the advantage that their consumer populations are large, which means you can bring new business model to scale quickly. Good examples here are Google or Amazon. Businesses in the 21st century will need to rethink the business model and business processes and most importantly get employees and consumers to work with systems and machines. For a company, the promise of AI technology is creating a win-win situation where it is making the humans super and not necessarily trying to make super humans.

In conclusion, the impacts and success factors across all three dimensions explore all primary corners of the organization which decisions makers can influence. If top management can not put together the right pods of skills, people and strategy can be extremely costly for the individual, top managing, the business and its end customers. The need to iterate strategies at an increased pace to keep a competitive advantages will take a toll on the company and they will need to learn how to deal with that. Management will need to make sure they invest in the right places and the right people are on board to push acceptance, adoption and adaption. Employees will see and feel change happening and it will be up to management to properly involve them. The objective of any AI project should be to have the buy-in from the target users and end users. Not doing so may lead to projects not getting the attention they deserve and end up failing not due to lack of feasibility but lazy or incentivized execution.

To make sure all future decision makers think about all the possible impacts and factors of integrating successfully, below will be a recommended checklist of question that must be asked regardless of the size of the company and the industry it is in.

7 Recommendation

Not all companies will have the experts or financial resources to get a jump start on AI, but that does not mean it's too late to participate in the AI race. With creative solutions and efficient planning, executives have the opportunity to make a difference and start adding long-term value to their companies with AI. Those in the position to design and implement technologies in their companies will still need to collect the human skills and capital to meet the final objective. In order to support decision makers to derive an actionable AI integration plan while being aware of the interrelatedness of impacts and complexity AI brings along, a recommendation in form of a checklist is outlined below. The checklist is a comprehensive list of questions based on all the topics discussed throughout this paper that should guide top managers towards the "best fit" integration of AI.

The increasing integration of AI application across industries is considered to be an accelerator for digital transformation and a large contributor to the disruption that will lead us into the fourth industrial revolution. Due to this, AI will have a profoundly transformative impact on our economy and society, for all of which we must prepare. Certain topics such as impact on law, policy and regulations will not be elaborated in detail, but the focus remains on the impacts and success factors of AI application integration. However, these topics are highly pertinent, which is why the next chapter briefly summarizes its importance and provides references for further research.

Since AI doesn't come with a "one-size-fits-all" technology, the first question which management needs to ask themselves is:

What's the best use of AI in your business?

Understanding the mission, vision and objective of the company reminds top-management what they are working for. If the purpose of the business is to provide a product or service or if quality, cost or speed are main drivers in the business model will be important to understand to find the best way to use AI and ML to add value towards that goal. For example, brands are finding use cases by offering relevant brand experiences like North Face (mission since 1996 is to provide best gear) working with IBM Watson

to match customers to the right gear or Sesame Street (mission is to help kids grow smarter, stronger and kinder at a local level) looking to deliver personalized learning experiences for children.

Under this umbrella, further questions to answer are:

- How can AI support your business's strategic objectives?
- What are your present and future business needs?
- Where could your business optimize processes with automation?
- What tasks could be automated to have an immediate benefit?
- Which of your competitors are applying AI and for what tasks?
- What success cases have other similar businesses created with AI solutions?

By answering these questions the business can quickly derive initial conclusions on whether there is potential for AI in the business, where the business stands in the Ai race in comparison to the rest of the competitors and potential uses cases the business can identify to show innovation and untapped growth potential.

The next key question to place is:

Can you start small and achieve organizational buy-in?

As mentioned before, AI and ML are not about automating jobs entirely or eliminating jobs and people. It is about augmenting the human capability to increase productivity and create downstream effects that can result in cost savings be it either from labor, clean data or increased efficiencies. In addition to a concrete benefit there should be an early involvement of the individuals that are key to the company's success Those people need to support the idea of AI integration from beginning and must believe in its benefits to act as "ambassadors" and therefore multipliers in the organization. Besides, these people will have the best ideas on the concrete use cases and quick wins of integrating AI. For example, a VP Support services might want to increase the quality of responses from support representatives or the Marketing Director might want to predict the acceptance of a new product based on sentiment from customer reviews and social media feeds. It will be key to start small and focused to not get overwhelmed by AIs complexity from the start.

Questions which can help to identify the use cases could be the following:

- What are you doing that you want to do faster, cheaper and better?
- What are you currently not doing that you want to start doing?
- What data is not available that could help make better decisions?
- What tasks are routine and add limited value?

By gathering information from various departments, executives may find patterns in outcomes. This can help to build a better business case and potentially creates the opportunity to accumulate budgets from different departments. Based on the list of cases with AI and ML potential, management can create a list of the roles which might be needed to be redefined in the future. By answering these questions, management prepares itself for challenging questions and for being able to convince the organization why a change towards AI is beneficial.

Once insights on potential use cases that support the business's objectives are there, the following question should be placed:

What is the status on your internal expertise?

As mentioned before, there is value in having the people who understand the business be part of the process. Their expertise can provide insights into challenges and opportunities when implementing a new technology. It is recommended to have a mix of people - on the one hand with technical focus, and on the other hand with business sense on the table. It can help to also actively involve people who openly express their doubts and concerns. They could take the role of challenging the applied methods. But, do you have these people at hand?

In order to identify the level of internal expertise and the respective consequences, following questions need to be answered:

- Is there someone who can lead an AI project from a technical and business needs perspective?
- Are there available data scientists or engineers with the skills to use ML or will you need to hire or consult someone?
- Do you have time and budget available to train your staff?
- If you want to outsource the development of your AI entirely to an external vendor, who is in charge of liaising with the vendor?

After answering these questions, we will understand if there is a talent gap or if the organization is well equipped to take on the challenge. As mentioned in the strategy section, it is important to have someone as a lead or a team with high level expertise available to provide guidance and see the project's needs in order for it to be successful.

In case of a decision for external support, following questions need to be added on top:

- What type of partnership is most suitable?
- Can the vendor scale with the project if it works?
- Shall we develop an in-house solution or use a third party one?
- What is the cost of experimenting in-house versus an external proof of concept?
- Is it possible to use open source software?

The answers to these questions will help to manage executive management expectation, to assess the ability to go to market and to mitigate the risk of not acting at all due to lack of internal know how or exclusivity with a vendor.

When having all of the organizational topics clear, it is time to focus on the "oil" of the AI machinery:

How good is your data?

- What kind of data is collected? Are there large volumes of structured or unstructured data?
- Does your core data provider maintain exclusive access to your company's data?
- Will your core data provider facilitate integration with your AI solution?
- Is the data properly structured for possible your AI's algorithms?
- If the data must be restructured, who will be responsible for this task?
- If you're missing any data you'll need, how will you obtain it?

These questions help understand the sources of data that are available for ML either it be structured or unstructured. Understanding this will help to identify the good from the more difficult ML problems to solve be it either predictive, for discovery or detection purposes.

The final questions in course of the AI integration should be around:

What is your AI implementation plan?

Using the information collected from the experts, understanding the use cases, the data available and having the organizational buy-in, it is time to sketch out the implementation plan.

- What is the desired timeline of the project?
- What are the target deadlines of major milestones such as pilot and launch?
- How will the project's ROI be measured?
- Which performance metrics will change as a result of implementing your AI solution?
- What tests will be used to measure these changes?
- What are the next steps in the AI strategy after implementing the project?
- How will you make sure to maintain and continuously improve your successfully implemented solution?

When running through these questions, the company might realize that it is not yet ready to implement and more information is needed. There might be also indications for a need to redesign some business processes along the way to make the implementation happen. The measurement of the project will be critical to the success of the project and a key indicator of its success to the people who use it and the decision makers who will provide the necessary support and funding to continue with more projects.

All of these questions demonstrate that the heavy work for AI needs to be done in the preparation phase and before the actual technical implementation can take place. The strategic considerations and the organizational readiness need to be addressed upfront. Although implementing AI is a big project it can and should be started in small with concrete use cases while always being aware of the big picture.

The AI integration process needs be driven with the respective management attention and can only succeed when decision makers manage to change the mindset and way of working for the company and its people.

8 Further research

The increasing integration of AI application across industries is will continue to be an accelerator for digital transformation and a large contributor to the disruption that will lead us into the fourth industrial revolution. Due to this, AI will have a profoundly transformative impact on our economy and society, all of which we must prepare. Further research areas not covered in this paper such as Law, Policy, and Regulation are highly pertinent. As mentioned these topics can pose great opportunity for challenges based on the types of laws, policies and regulations companies might have to abide to in their respective countries of operations. All top management should be aware of the trends in these areas and how they might affect short or long term strategies. Below are short summaries on their importance and references to channels for further research.

8.1 Law

As companies build, implement and integrate AI to automate tasks, augment the human decision capabilities or add autonomy, there will be new areas of concern related to law and how it complies with it. Jerry Kaplan, an artificial intelligence expert, discusses this topic in his book “Artificial Intelligence: What everyone needs to know” by identifying some questions we as a society and governing judicial bodies should consider, such as (Kaplan, 2015):

- Can a computer program enter in agreement and contracts?
- Should an intelligent agent be limited in what it is permitted to do?
- Should people take on full responsibility for their intelligent agent?
- Can an AI system commit a crime?
- Can an AI system be held accountable for a criminal crime?

The discussions to find answers to these questions have begun. Companies that have already integrated AI into their strategy will have many people watching on how their AI systems are held accountable. They will be the first to be set as examples. Tesla,

an electric car manufacturer, built a semi-autonomous autopilot system into the Tesla Model S. While in autopilot, the vehicle slammed into the side of a tractor trailer turning across his path and the driver became the first person to die in a partially autonomous car (Stewart, 2017). The accident was outside the capabilities of the Autopilot and Automatic Emergency Braking (AEB) system and since Tesla notified the driver that the autopilot system needed the driver's supervision, it was not found guilty of having faulty hardware or software.

The investigation by the NHTSA was extensive and detailed which helped Tesla update their existing software. Although Tesla was not at fault, there are still over 20 other automakers who are also working on introducing driver assisted technologies which might come across the same situation and can learn from them. As governments catch up to the technological development, companies should keep it a priority to do their due diligence on the changes in AI law that might affect them if they were to be in a situation where their AI system was to be part of a crime, accident or malpractice.

There are a couple of associations, and references companies can refer to stay up-to-date such as:

- International Association for Artificial Intelligence and Law (IAAIL)
- The Foundation for Legal Knowledge Based Systems (JURIX)
- Artificial Intelligence and Law Journal

8.2 Policy and regulation

According to the Stanfords One Hundred Year Study on Artificial Intelligence (AI100), public policy should help society adapt to AI applications to reap its benefits, mitigate their inevitable errors and overcome failures (Stanford, 2016). According to the AI100 Study Panel from Stanford, they recommend all layers of government acquire technical expertise in AI.

AI applications have the potential and are on course to amplifying gap of opportunity between those who have and do not have access. Due to this, policies should make sure to be evaluated on criteria such as enabling democratic values, sharing of AI's benefits, or avoiding only concentrated power and profits in the hands of a fortunate few. The companies who participate in the AI revolution will be significant influences on such policies. It will be their corporate social responsibility to understand the importance

of these criteria and take on a role that contributes back to society in positive ways and reduce negative impact.

An important aspect of policy development in AI is regulation. Policy combined with the regulatory approach must adhere to three main elements: compliance, multi-stakeholder approach, and international cooperation(Gummi, 2017). Regulation can either push AI forward or hinder AI progress. It can push forward in new areas previously not actively regulated, or it can delay where the law is strict and complying with it is difficult due to its strictness . AI driven companies, governments and especially the people will need to work very closely together to regulate positively. Governments have already started discussions on this policy and regulations to inform the general public what to expect and what challenges need to be resolved.

Below are publications from governments who have acknowledged the coming of AI and its impact on regulation:

- The United States White House: Artificial Intelligence, Automation, and the Economy (Published December 2016)
- European Parliament Legal Affairs Committee: European Civil Law Rules in Robotics (Published October 2016)
- House of Commons, United Kingdom: Robotics and artificial intelligence (Published September 2016)

9 Conclusion

To understand the oncoming impacts and success factors to successfully integrate AI technologies will only be half the battle for most companies who find themselves looking for ways to get into the world of AI. The impacts on organizations are real and successful integration will be challenging, but with the right mindset and people to drive the initiative forward, integrating AI will be the best investment since the internet.

A key take away from this paper is to understand that AI's biggest problem is that it is artificial. The technologies are artificially created through software developers algorithms, fed with data and powered by cloud-based servers - all controlled by humans. Humans are the most important success factor on many levels, such as making sure the machine is doing what it is supposed to be doing, that the information it is displaying makes sense for its ultimate outcome, and that it is reaching its outcome in the way it is meant to be behaving. Many argue that overlooking this human factor can be dangerous and thus is critical for a successful AI integration.

All companies - them being in the right industry or sector, having lots of talent or close to none, having started on implementing AI or not knowing where to start - will need to answer the recommended set of questions in one form or another. From machine learning to deep learning technologies, the algorithms behind all the "learning" will be critical to a company's success in choosing or tweaking the right one to make great predictions and do it more and more accurately as they learn from past experiences. However, they also have risks associated to them. These risks are most common when we do not understand either the algorithm, the input data it is crunching or the environment it's making predictions for. Informing ourselves about the capabilities and limitations will be essential in setting realistic expectations and scaling for the long-term. Yet, we will also need to start becoming comfortable with not knowing the exact "how" since up today the way machines learn is still a black box. Those who will benefit the most will be the ones who understand the algorithms and have access to data.

Through leadership, a thorough strategy that's inclusive of its employees and observant of its real needs, companies will have higher chances of succeeding in its AI endeavors. Communication and knowledge sharing management will be a key success factor to integrate AI. Without the buy-in of the key stakeholders in the organization, the integration of AI can't become a success story.

List of Figures

4.1	Datasets Over Algorithms	24
4.2	Companies explore different ways to Develop AI Capabilities	27
5.1	Supervised Learning	36
5.2	Unsupervised Learning	38
5.3	Deep learning	40
6.1	Sectors Adoption	47
6.2	Companies Explore Different Ways to Develop AI Capabilities	66

References

- Anthony, S. (2012). *First mover or fast follower?* Harvard Business Review.
- Benedikt F. C., M., Osborne. (2013). The future of employment.
- BI, I. (2017). *4relations*. Business Insider. Retrieved from <http://www.businessinsider.de/chatbots-vs-humans-for-customer-relations-2016-12?r=US&IR=T>
- Bjorn, C. (2016). Smartphone speech recognition can write text messages three times faster than human typing. Retrieved from <http://news.stanford.edu/2016/08/24/stanford-study-speech-recognition-faster-texting/>
- Brokaw, A. (2016). *This startup uses machine learning and satellite imagery to predict crop yields*. Retrieved from <https://www.theverge.com/2016/8/4/12369494/descartes-artificial-intelligence-crop-predictions-usda>
- Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies* (1st ed.). W. W. Norton & Company.
- Brynjolfsson, E., & McAfee, A. (2017). *Whats driving the machine learning explosion?* Retrieved from <https://hbr.org/2017/07/whats-driving-the-machine-learning-explosion+&cd=5&hl=en&ct=clnk&gl=us>
- Buck, I. (2015). *NVIDIA's next-gen pascal GPU architecture to provide 10x speedup for deep learning apps*. Retrieved from <https://blogs.nvidia.com/blog/2015/03/17/pascal/>
- Chui, M. a. (2015). Four fundamentals of workplace automation. *Mckinsey*.
- Columbus, L. (2017). *Mckinsey's state of machine learning and ai*. Forbes. Retrieved from <https://www.forbes.com/sites/louiscolombus/2017/07/09/mckinseys-state-of-machine-learning-and-ai-2017/#43f5e9cf75b6>
- Davenport, T. H., & Kirby, J. (2016). *Only humans need apply: Winners and losers in the age of smart machines*. NY: HarperCollins.
- Denning, P. J., & Lewis, T. G. (2016). Exponential laws of computing growth. *Commun. ACM*, 60(1), 54-65. doi: 10.1145/2976758
- Dorrier, J. (2016). *Will the end of moore's law halt computing's exponential rise?* Retrieved from <https://singularityhub.com/2016/03/08/will-the-end-of-moores-law-halt-computings-exponential-rise/>
- Esteva, A. e. a. (2017). *Dermatologist-level classification of skin cancer with deep neural networks*. Nature.
- Fast, E., & Horvitz, E. (2017). Long-term trends in the public perception of artificial intelligence.

-
- Gaines-Ross, L. (2017, Apr). *What do people - not techies, not companies - think about artificial intelligence?* Retrieved from <https://hbr.org/2016/10/what-do-people-not-techies-not-companies-think-about-artificial-intelligence>
- Gownder, J. (2017). TechradarTM: Automation technologies, robotics, and ai in the workforce, q2 2017. *Forrester*.
- Guardian, T. (2014). *Elon musk: Artificial intelligence is our biggest existential threat*. Retrieved from <https://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat>
- Gummi, M. (2017). *Artificial intelligence:the way forward for policy and regulation*. Retrieved from <https://berkeleypublicpolicyjournal.org/2017/04/12/artificial-intelligence-the-way-forward-for-policy-and-regulation/>
- Harris, D. (2015). *How paypal uses deep learning and detective work to fight fraud*. Retrieved from https://gigaom.com/2015/03/06/how-paypal-uses-deep-learning-and-detective-work-to-fight-fraud/?utm_content=buffer9cf7e&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer
- Iansiti, M., & West, J. (1997). *Technology integration:turning great research into great products*. Retrieved from <https://webcache.googleusercontent.com/search?q=cache:tp8tUisJ3cEJ:https://hbr.org/1997/05/technology-integration-turning-great-research-into-great-products+&cd=1&hl=en&ct=clnk&gl=us>
- Ito, J., & Howe, J. (2016). *Whiplash: How to survive our faster future*. Grand Central Publishing.
- Jia, Y., & Vajda, P. (2016). *Delivering real-time ai in the palm of your hand*. Retrieved from <https://code.facebook.com/posts/196146247499076/delivering-real-time-ai-in-the-palm-of-your-hand/>
- Kaplan, J. (2015). *Humans need not apply: A guide to wealth and work in the age of artificial intelligence*. NH: Yale University Press.
- Kurzweil, R. (2001). *The law of accelerating returns*. Retrieved from <http://www.kurzweilai.net/the-law-of-accelerating-returns>
- Lardinois, F. (2014). *Google announces massive price drops for its cloud computing services and storage,introduces sustained-used discounts*. Retrieved from <https://techcrunch.com/2014/03/25/google-drops-prices-for-compute-and-app-engine-by-over-30-cloud-storage-by-68-introduces-sustained-use-discounts/>

- Markoff, J. (2014). The competitive landscape for machine intelligence. *New York Times*. Retrieved from <https://www.nytimes.com/2014/12/16/science/century-long-study-will-examine-effects-of-artificial-intelligence.html>
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (1955). *A proposal for the dartmouth summer research project on artificial intelligence*. Retrieved from <http://www-formal.standford.edu/jmc/history/dartmouth/dartmouth.html>
- Mitchell, T. M. (1997). *Machine learning*. McGraw Hill.
- News, C. (2016). 60 minutes/vanity fair poll: Artificial intelligence. Retrieved from <https://www.cbsnews.com/news/60-minutes-vanity-fair-poll-artificial-intelligence/>
- Ng, A. (2016). *Hiring your first chief ai officer*. Retrieved from <https://hbr.org/2016/11/hiring-your-first-chief-ai-officer>
- Nilsson, N. J. (2010). *The quest for artificial intelligence: A history of ideas and achievements*. NY: Cambridge University Press.
- Petropoulos, G. (2017). *Do we understand the impact of artificial intelligence on employment?* Retrieved from <http://bruegel.org/2017/04/do-we-understand-the-impact-of-artificial-intelligence-on-employment/>
- Purdy, M., & Daugherty, P. (2016). *Why artificial intelligence is the future of growth*. Retrieved from https://www.accenture.com/t20170206T005353_w_us-en/_acnmedia/PDF-33/Accenture-Why-AI-is-the-Future-of-Growth.PDF#zoom=50
- Russell, S. J., & Norvig, P. (2003). *Artificial intelligence: A modern approach* (2nd ed.). Pearson Education.
- Samuel, A. L. (1959). Some studies in machine learning using the game of checkers. *IBM Journal*, 3(3), 210-229.
- Sculley, D. a. (2015). Hidden technical debt in machine learning systems. In *Proceedings of the 28th international conference on neural information processing systems* (pp. 2503–2511). Cambridge, MA, USA: MIT Press. Retrieved from <http://dl.acm.org/citation.cfm?id=2969442.2969519>
- Sharma, V. (2016). Malaria outbreak prediction model using machine learning. Retrieved from <http://ijarcet.org/wp-content/uploads/IJARCET-VOL-4-ISSUE-12-4415-4419.pdf>
- Stewart, J. (2017). *After probing teslas deadly crash, feds say yay to self-driving*. Retrieved from <https://www.wired.com/2017/01/probing-teslas-deadly-crash-feds-say-yay-self-driving/>

- Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., & Hager, G. (2016). *Artificial intelligence and life in 2030*. Retrieved from https://ai100.stanford.edu/sites/default/files/ai_100_report_0831fnl.pdf
- Thompson, B. (2017). The arrival of artificial intelligence. *Stratechery*. Retrieved from <https://stratechery.com/2017/the-arrival-of-artificial-intelligence/>
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433-460.
- Westerman, G., Bonnet, D., & McAfee, A. (2014). Leading digital: Turning technology into business transformation. *Harvard Business Review*.
- Westermann, G., Tannou, M., Bonnet, D., Ferraris, P., & McAfee, A. (2012). *The digital advantage: How digital leaders outperform their peers in every industry*. Retrieved from <https://www.capgemini.com/resources/the-digital-advantage-how-digital-leaders-outperform-their-peers-in-every-industry>
- Wissner Gross, A. (2016). Datasets over algorithms. Retrieved from <https://www.edge.org/response-detail/26587>
- Yeomans, M. (2015). What every manager should know about machine learning. *Harvard Business Review*. Retrieved from <https://hbr.org/2015/07/what-every-manager-should-know-about-machine-learning>