

***NEAPOLIS UNIVERSITY OF
PAPHOS (NUP)***

WORKING PAPERS SERIES

2017/01

TITLE: The Forthcoming Artificial Intelligence Revolution

AUTHOR: Spyros Makridakis

The Forthcoming Artificial Intelligence (AI) Revolution: Its Impact on Society and Firms

Spyros Makridakis

Rector of Neapolis University Pafos

The rise of powerful AI will be either the best or the worst thing ever to happen to humanity. We do not yet know which.

Stephen Hawking

Abstract

The impact of the industrial and digital (information) revolutions has, undoubtedly, been substantial on practically all aspects of our society, life, firms and employment. Will the forthcoming AI revolution produce similar, far-reaching effects? By examining analogous inventions of the industrial, digital and AI revolutions, this article claims that the latter is on target and that it would bring extensive changes that will also affect all aspects of our society and life. In addition, its impact on firms and employment will be considerable, resulting in richly interconnected organizations with decision making based on the analysis and exploitation of “big” data and intensified, global competition among firms. People will be capable of buying good and obtaining services from anywhere in the world using the Internet, and exploiting the unlimited, additional benefits that will open through the widespread usage of AI inventions. The paper concludes that significant competitive advantages will continue to accrue to those utilizing the Internet widely and willing to take entrepreneurial risks in order to turn innovative products/services into worldwide commercial success stories. The greatest challenge facing societies and firms would be utilising the benefits of availing AI technologies, providing vast opportunities for both new products/services and immense productivity improvements while avoiding the dangers and disadvantages in terms of increased unemployment and greater wealth inequalities.

KEYWORDS: Artificial Intelligence (AI); Industrial Revolution; Digital Revolution; AI Revolution; Impact of AI Revolution; Benefits and Dangers of AI Technologies.

Over the past decade, numerous predictions have been made about the forthcoming Artificial Intelligence (AI) Revolution and its impact on all aspects of our society, firms and life in general. This paper considers such predictions and compares them to those of the industrial and digital ones. A similar paper was written by this author and published in this journal in 1995, envisioning the forthcoming changes being brought by the digital (information) revolution, developing steadily at that time, and predicting its impact for the year 2015 (Makridakis, 1995). The current paper evaluates these 1995 predictions and their impact identifying hits and misses with the purpose of focusing on the new ones being

brought by the AI revolution. It must be emphasized that the stakes of rightly predicting the impact of the AI revolution are far reaching as intelligent machines may become our “final invention” that may end human supremacy (Barat, 2013). There is little doubt that AI holds enormous potential as computers and robots will probably achieve, or come close to, human intelligence over the next twenty years becoming a serious competitor to all the jobs currently performed by humans and for the first time raising doubt over the end of human supremacy.

This paper is organized into four parts. It first overviews the predictions made in the 1995 paper for the year 2015, identifying successes and failures and concluding that major technological developments (notably the Internet and smartphones) were undervalued while the general trend leading up to them was predicted correctly. Second, it investigates existing and forthcoming technological advances in the field of AI and the ability of computers/machines to acquire real intelligence. Moreover, it summarizes prevailing, major views of how AI may revolutionize practically everything and its impact on the future of humanity. The third section sums up the impact of the AI revolution and describes the four major scenarios being advocated, as well as what could be done to avoid the possible negative consequences of AI technologies. The fourth section discusses how firms will be affected by these technologies that will transform the competitive landscape, how start-up firms are founded and the way success could be achieved. Finally, there is a brief concluding section speculating about the future of AI and its impact on our society, life, firms and employment.

Table 1 summarizes the major inventions of the industrial, digital and AI revolutions. The first two columns have been taken from the 1995 paper (Makridakis, 1995) replacing “Widespread use of” with “Actual use in 2015” while the third one is new referring to the AI revolution and its existing and new, widespread used inventions by the year 2037. The challenge is to correctly predict the impact of AI inventions and how the role of humans will be affected when machines of equal, or superior intelligence could substitute, supplement and/or amplify practically all mental tasks that until now have been the exclusive province of humans. As with technological predictions made in the past, as well as those in the 1995 paper, it is necessary not to extrapolate in a linear fashion in order not to underestimate the rate of technological change and its impact on all aspects of our society, life, work and firms. In my opinion the forthcoming technologies of the AI revolution and their impact over the next twenty years will probably be many times the magnitude of those of the digital one from 1995 to 2015 and probably even greater than those of the Industrial revolution.

Table 1: From Steam Engines to Unattended Factories, from the ENIAC Computer to Big Data and from Neural Net Devices to Self-Driving Cars and Singularity

Industrial Revolution

(Mechanical power)

Substituting, supplementing and/or amplifying routine manual tasks

1712 Newcomen's steam engine
 1784 Watt's double action steam engine
 1830 Electricity
 1876 Otto's internal combustion engine
 1890 Cars
 1901 Electricity in homes
 1914 Continuous production line

 1919 Electricity in one-third of homes

Actual use of:

1950s Electrical appliances
 1960s Cars

 1970s Long-distance telephones

 2010 Unattended factories

Digital Revolution

(Computer power)

Substituting, supplementing and/or amplifying standardized mental tasks

1946 ENIAC Computer
 1950s IBM's business computers
 1970s Electronic data processing (EDP)
 1971 Time-sharing computers
 1973 Microprocessor
 1977 Apple's computer
 1980s Computers with modems

 1993 Personal computers in one-third of homes

Actual use in 2015:

2015 61% of people use smartphones
 2015 Amazon most valuable US retailer
 (Surpassing Walmart)
 2015 23% of employees work from home
 (Full-time or part-time)
 2015 Collecting/Exploiting Big Data

AI Revolution

(Brain power)

Substituting, supplementing and/or amplifying practically all mental tasks

1990 Neural net device reads handwritten digits
 1993 Robot Polly navigates using vision
 1997 Deep Blue defeats the world chess champion
 1998 Robotic toy Furby learns how to speak
 2005 Robot ASIMO serves restaurant customers
 2009 Google's first self-driving car
 2011 Watson computer beats Jeopardy's
 champions
 2016 AlphaGo defeats GO champions using
 neural learning algorithms

Widespread use of:

202? Computer translations
 202? Self-driving cars

 202? Deep neural learning

 203? Machines reach human intelligence

The 1995 Paper: Hits and Misses

The 1995 paper (Makridakis, 1995) was written at a time when the digital (at that time it was called information) revolution was progressing at a steady rate. The paper predicted that by 2015 *“the information revolution should be in full swing”* and that *“computers/communications”* would be in widespread use, which has actually happened, although its two most important inventions (the Internet and smartphones) and their significant influence were not foreseen as such. Moreover, the paper predicted that *“a single computer (but not a smartphone) can, in addition to its traditional tasks, also become a terminal capable of being used interactively for the following:”* (p. 804-805)

- Picturephone and teleconference
- Television and videos
- Music
- Shopping
- On line banking and financial services
- Reservations
- Medical advice
- Access to all types of services
- Video games
- Other games (e.g., gambling, chess etc.)
- News, sports and weather reports
- Access to data banks

The above have all materialized and can indeed be accessed by computer, although the extent of their utilization was underestimated as smartphones are now being used widely. For instance, the ease of accessing and downloading scientific articles on one’s computer in his/her office or home would have seemed like science fiction back in 1995, when finding such articles required spending many hours in the library (often in its basement for older publications) and making photocopies to keep them for later use. Moreover, having access, from one’s smartphone or tablet, to news from anywhere in the world, being able to subscribe to digital services, obtain weather forecasts, purchase games, watch movies, make payments using smartphones and a plethora of other, useful applications was greatly underestimated, while the extensive use of the cloud for storing large amounts of data for free was not predicted at all at that time. Even in 1995 when the implications of Moore’s law leading to increasing computer speed and storage while reducing costs were well known, nevertheless, it was hard to imagine that in 2016 there would be 60 trillion web pages, 2.5 billion smartphones, more than 2 billion personal computers and 3.5 billion Google searches a day.

The paper correctly predicted *“as wireless telecommunications will be possible the above list of capabilities can be accessed from anywhere in the world without the need for regular telephone lines”*. What the 1995 paper missed, however, was that in 2015 top smartphones, costing less than €500, would be as powerful as the 1995 supercomputer, allowing access to the Internet and all tasks that were only performed by expensive computers at that time, including an almost unlimited availability of new, powerful apps providing a large array of innovative services that were not imagined twenty years ago. Furthermore, the paper correctly predicted super automation leading to unattended factories stating that *“by 2015*

there will be little need for people to do repetitive manual or mental tasks". It also foresaw the decline of large industrial firms, increased global competition and the drop in the percentage of labour force employed in agriculture and manufacturing (more on these predictions in the section **The Impact of the AI Revolution on Firms**). It missed however the widespread utilization of the Internet (at that time it was a text only service), as well as search engines (notably Google), social networking sites (notably Facebook) and the fundamental changes being brought by the widespread use of Apple's iPhone, Samsung's Galaxy and Google's Android smartphones. It is indeed surprising today to see groups of people in a coffee shop or restaurant using their smartphones instead of speaking to each other and young children as little as three or four years of age playing with phones and tablets. Smartphones and tablets connected to the Internet through Wi-Fi have influenced social interactions to a significant extent, as well as the way we search for information, use maps and GPS for finding locations, as well as making payments. These technologies were not predicted in the 1995 paper.

Towards the AI Revolution

The 1995 paper referred to Say, the famous French economist, who wrote in 1828 about the possibility of cars as substitutes for horses:

"Nevertheless no machine will ever be able to perform what even the worst horses can—the service of carrying people and goods through the bustle and throng of a great city." (p. 800)

Say could have never dreamed of, in his wildest imagination, self-driving cars, pilotless airplanes, Skype calls, super computers, smartphones or intelligent robots. Technologies that seemed like pure science fiction less than 190 years ago are available today and some like self-driving vehicles will in all likelihood be in widespread use within the next twenty years. The challenge is to realistically predict forthcoming AI technologies without falling into the same short-sighted trap of Say and others, including my 1995 paper, unable to realize the momentous, non-linear advancements of new technologies. There are two observations to be made.

First, 190 years is a brief period by historical standards and during this period we went from horses being the major source of transportation to self-driving cars and from the abacus and slide rules to powerful super computers in our pockets. Secondly, the length of time between technological inventions and their practical, widespread use is constantly being reduced. For instance, it took more than 200 years from the time Newcomen developed the first workable steam engine in 1707 to when Henry Ford built a reliable and affordable car in 1908. It took more than 90 years between the time electricity was invented and its extensive use by firms to substantially improve factory productivity. It took twenty years, however, between ENIAC, the first computer, and IBM's 360 system that was mass produced and was affordable by smaller business firms while it took only ten years between the invention of the mobile phone in 1973 by Dr Martin Cooper and its public launch by Motorola. The biggest and most rapid progress, however, took place with smartphones which first appeared in 2002 and saw a stellar growth with the release of new versions possessing substantial improvements every one or two years by the likes of Apple, Samsung and several Chinese firms. Smartphones, in addition to their technical features, now incorporate artificial intelligence characteristics that include understanding speech, providing customized advice in spoken language, completing words when writing a text and

several other functions requiring embedded AI, provided by a pocket computer smaller in size than a pack of cigarettes.

From smart machines to clever computers and to Artificial Intelligence (AI) programs: A thermostat is a simple mechanical device exhibiting some primitive but extremely valuable type of intelligence by keeping temperatures constant at some desired, pre-set level. Computers are also clever as they can be instructed to make extremely complicated decisions taking into account a large number of factors and selection criteria, but like thermostats such decisions are pre-programmed and based on logic, if-then rules and decision trees that produce the ***exact*** same results, as long as the input instructions are alike. The major advantage of computers is their lightning speed that allows them to perform billions of instructions per second. AI, on the other hand, goes a step further by not simply applying pre-programmed decisions, but instead exhibiting some learning capabilities.

The reading of handwritten digits (first utilized to determine the written amount on bank checks) by the neural net device in 1990 (see Table 1) is a predecessor of efforts to achieve learning. Handwritten digits could be inscribed in almost infinite ways so programming a machine to correctly read them was no small challenge, beginning a new direction in programming aimed at mimicking the human mind that can effortlessly understand all types of handwritten digits. The same applies for Polly (see Table 1) that could learn to pronounce a number of words, or ASIMO that could navigate in open places of constantly changing environmental settings, rendering preprogramming impossible to account for all possible cases. ASIMO's decisions were quite different than those of the extremely fast supercomputer Deep Blue (see Table 1) that used "brute force" to identify and analyze up to 60 billion moves within the three minute period allowed to chess players to make their next move. Deep Blue was incapable of learning as its programming was based on logic, if-then rules and decision trees which meant that it could not learn and could make the same mistakes over and over again.

The story of the Watson computer beating Jeopardy's two most successful contestants is more complicated, since retrieving the most appropriate answer out of the 200 million pages of information stored in its memory is not a sign of real intelligence as it relied on its lightning speed to retrieve information in seconds. What is more challenging according to Jennings, one of Jeopardy's previous champions, is *"to read clues in a natural language, understand puns and the red herrings, to unpack just the meaning of the clue"* (May, 2013). Similarly, it is a sign of intelligence to improve its performance by *"playing 100 games against past winners"*. (Best, 2016). Watson went several steps beyond Deep Blue towards AI by being able to understand spoken English and learn from his mistakes (New Yorker, 2016). However, he was still short of AlphaGo that defeated Go Champions in a game that cannot be won simply by using "brute force" as the number of moves in this game is infinite, requiring the program to use learning algorithms that can improve its performance as it plays more and more games

Computers and real learning: According to its proponents, *"the main focus of AI research is in teaching computers to think for themselves and improvise solutions to common problems"* (Clark, 2015). But many doubt that computers can learn to think for themselves even though

they can display signs of intelligence. David Silver, an AI scientist working in DeepMind, explained that *“even though AlphaGo has affectively rediscovered the most subtle concepts of Go, its knowledge is implicit. The computer parse out these concepts – they simply emerge from its statistical comparisons of types of winning board positions at GO”* (Chouard, 2016). At the same time Cho Hyeon, one of the strongest Go players in Korea commented that *“AlphaGo seems like it knows everything!”* while others believe that *“AlphaGo is likely to start a ‘new revolution’ in the way we play Go”* as *“it is seeking simply to maximize its probability of reaching winning positions, rather than as human players tend to do – maximize territorial gains”* (Chouard, 2016). Does it matter, as Silver said, that AlphaGo’s knowledge of the game is implicit as long as it can beat the best players? A more serious issue is whether or not AlphaGo’s ability to win games with fixed rules can extend to real life settings where not only the rules are not fixed, but they can change with time, or from one situation to another.

From digital computers to AI tools: The Intel Pentium microprocessor, introduced in 1993, incorporated graphics and music capabilities and opened computers up to a large number of affordable applications extending beyond just data processing. Such technologies signalled the beginning of a new era that now includes intelligent personal assistants understanding and answering natural languages, robots able to see and perform a intelligent functions, self-driving vehicles and a host of other capabilities which were until then an exclusive human ability. The tech optimists ascertain that in less than 25 years computers went from just manipulating 0 and 1 digits, to utilizing sophisticated neural network algorithms that enable vision and the understanding and speaking of natural languages among others. Technology optimists therefore maintain there is little doubt that in the next twenty years, accelerated AI technological progress will lead to a breakthrough, based on deep learning that imitates the way young children learn, rather than the laborious instructions by tailor-made programs aimed for specific applications and based on logic, if-then rules and decision trees (Parloff, 2016).

For instance, DeepMind is based on a neural program utilizing deep learning that teaches **itself** how to play dozens of Atari games, such as Breakout, as well or better than humans, without specific instructions for doing so, but by playing thousands of games and **improving** itself each time. This program, trained in a different way, became the AlphaGo that defeated GO champion Lee Sodol in 2016. Moreover, it will form the core of a new project to learn to play Starcraft, a complicated game based on both long term strategy as well as quick tactical decisions to stay ahead of an opponent, which DeepMind plans to be its next target for advancing deep learning (Kahn, 2016). Deep learning is an area that seems to be at the forefront of research and funding efforts to improve AI, as its successes have sparked a burst of activity in equity funding that reached an all-time high of more than \$1 billion with 121 projects for start-ups in the second quarter of 2016, compared to 21 in the equivalent quarter of 2011 (Parloff, 2016).

Google had two deep learning projects underway in 2012. Today it is pursuing more than 1,000, according to their spokesperson, in all its major product sectors, including search, Android, Gmail, translation, maps, YouTube, and self-driving cars (The Week, 2016). IBM’s Watson system used AI, but not deep learning, when it beat the two Jeopardy champions in 2011. Now though, almost all of Watson’s 30 component

services have been augmented by deep learning. Venture capitalists, who did not even know what deep learning was five years ago, today are wary of start-ups that do not incorporate it into their programs. We are now living in an age when it has become mandatory for people building sophisticated software applications to avoid click through menus by incorporating natural-language processing tapping deep learning (Parloff, 2016).

How far can deep learning go? There are no limits according to technology optimists for two reasons. First as progress is available to practically everyone to utilize through *Open Source* software, researchers will concentrate their efforts on new, more powerful algorithms leading to cumulative learning. Second and equally important, in the future intelligent computer programs will be capable of writing new programs themselves, initially perhaps not so sophisticated ones, but improving with time as learning will be incorporated to be part of their abilities. Kurzweil (Kurzweil, 2005) sees nonbiological intelligence to match the range and subtlety of human intelligence within a quarter of a century and what he calls “Singularity” to occur by 2045, bringing *“the dawning of a new civilization that will enable us to transcend our biological limitations and amplify our creativity. In this new world, there will be no clear distinction between human and machine, real reality and virtual reality”*. For some people these predictions are startling, with far-fetched implications should they come true. In the next section, four scenarios associated with the AI revolution are presented and their impact on our societies, life work and firms is discussed.

The Four AI Scenarios

Until rather recently, famines, wars and pandemics were common, affecting sizable segments of the population, causing misery and devastation as well as a large number of deaths. The industrial revolution considerably increased the standards of living while the digital one maintained such rise and also shifted employment patterns, resulting in more interesting and comfortable office jobs. The AI revolution is promising even greater improvements in productivity and further expansion in wealth. Today more and more people, at least in developed countries, die from overeating rather than famine, commit suicide instead of being killed by soldiers, terrorists and criminals combined and die from old age rather than infectious disease (Harari, 2016). Table 1 shows the power of each revolution with the industrial one aiming at routine manual tasks, the digital doing so to routine mental ones and AI aiming at substituting, supplementing and/or amplifying practically **all** tasks performed by humans. The critical question is: *“what will the role of humans be at a time when computers and robots could perform as well or better **and** much cheaper, practically all tasks that humans do at present?”* There are four scenarios attempting to answer this question.

The Optimists: Kurzweil and other optimists predict a “science fiction”, utopian future with Genetics, Nanotechnology and Robotics (GNR) revolutionizing everything, allowing humans to harness the speed, memory capacities and knowledge sharing ability of computers and our brain being directly connected to the cloud. Genetics would enable changing our genes to avoid disease and slow down, or even reverse aging, thus extending our life span considerably and perhaps eventually achieving immortality. Nanotechnology would enable us to create virtually any physical product from information and inexpensive materials bringing an unlimited creation of wealth. Finally, robots would be doing all the actual work,

leaving humans with the choice of spending their time performing activities of their choice and working, when they want, at jobs that interest them.

The Pessimists: In a much quoted article from *Wired* magazine in 2000, Bill Joy (Joy, 2000) wrote “*Our most powerful 21st-century technologies – robotics, genetic engineering, and nanotech – are threatening to make humans an endangered species*”. Joy pointed out that as machines become more and more intelligent and as societal problems become more and more complex, people will let machines make all the important decisions for them as these decisions will bring better results than those made by humans. This situation will, eventually, result in machines being in effective control of all important decisions with people dependent on them and afraid to make their own choices. Joy and many other scientists (Cellan-Jones, 2016) and philosophers (Bostrom, 2014) believe that Kurzweil and his supporters vastly underestimate the magnitude of the challenge and the potential dangers which can arise from thinking machines and intelligent robots. They point out that in the utopian world of abundance, where all work will be done by machines and robots, humans may be reduced to second rate status (some saying the equivalent of computer pets) as smarter than them computers and robots will be available in large numbers and people will not be motivated to work, leaving computers/robots to be in charge of making all important decisions. It may not be a bad world, but it will be a different one with people delegated to second rate status.

Harari is the newest arrival to the ranks of pessimists. His recent book (Harari, 2016, p. 397) concludes with the following three statements:

- “Science is converging to an all-encompassing dogma, which says that organisms are algorithms, and life is data processing”
- “Intelligence is decoupling from consciousness”
- “Non-conscious but highly intelligent algorithms may soon know us better than we know ourselves”

Consequently, he asks three key questions (which are actually answered by the above three statements) with terrifying implications for the future of humanity:

- “Are organisms really just algorithms, and is life just data processing”
- “What is more valuable – intelligence or consciousness?”
- “What will happen to society, politics and daily life when non-conscious but highly intelligent algorithms know us better than we know ourselves?”

Harari admits that nobody really knows how technology will evolve or what its impact will be. Instead he discusses the implications of each of his three questions:

- If indeed organisms are algorithms then thinking machines utilizing more efficient ones than those by humans will have an advantage. Moreover, if life is just data processing then there is no way to compete with computers that can consult/exploit practically all available information to base their decisions.
- The non-conscious algorithms Google search is based on the consultation of millions of possible entries and often surprise us by their correct recommendations,. The implications that similar, more advanced algorithms than those utilized by Google search will be developed (bearing in mind Google search is less than twenty years old) in the future and be able to access all available information from complete data

bases are far reaching and will “provide us with better information than we could expect to find ourselves”.

- Humans are proud of their consciousness, but does it matter that self-driving vehicles do not have one, but still make better decisions than human drivers, as can be confirmed by their significantly lower number of traffic accidents?

When AI technologies are further advanced and self-driving vehicles are in widespread use, there will come a time that legislation will be passed forbidding human driving. Clearly, self-driving vehicles do not exceed speed limits, do not drive under the influence of alcohol or drugs, do not get tired, do not get distracted by talking on the phone or sending emails and in general make fewer mistakes than human drivers, causing fewer accidents. There are two implications if humans are not allowed to drive. First, there will be a huge labour displacement for the 3.5 million unionized truck drivers in the USA and the 600 thousand ones in the UK (plus the additional number of non-unionized ones) as well as the more than one million taxi and Uber drivers in these two countries. Second, and more importantly, it will take away our freedom of driving, admitting that computers are superior to us. Once such an admission is accepted there will be no limits to letting computers also make a great number of other decisions, like being in charge of nuclear plants, setting public policies or deciding on optimal economic strategies as their biggest advantage is their objectivity and their ability to make fewer mistakes than humans.

One can go as far as suggesting letting computers choose Presidents/Prime Ministers and elected officials using objective criteria rather than having people voting emotionally and believing the unrealistic promises that candidates make. Although such a suggestion will never be accepted, at least not in the near future, it has its merits since people often choose the wrong candidate and later regret their choice after finding out that pre-election promises were not only broken, but they were even reversed. Critics say if computers do eventually become in charge of making all important decisions there will be little left for people to do as they will be demoted to simply observing the decisions made by computers, the same way as being a passenger in a car driven by a computer, not allowed to take control out of the fear of causing an accident. As mentioned before, this could lead to humans eventually becoming computers’ pets.

The Pragmatists: At present the vast majority of views about the future implications of AI are negative, concerned with its potential dystopian consequences (Elon Musk, the CEO of Tesla, says it is like “summoning the demon” and calls the consequences worse than what nuclear weapons can do). There are fewer optimists and only a couple of pragmatists like Sam Altman and Michio Kaku (Peckham, 2016) who believe that AI technologies can be controlled through “OpenAI” and effective regulation. The ranks of pragmatists also includes John Markoff (Markoff, 2016) who pointed out that the AI field can be distinguished by two categories: The first trying to duplicate human intelligence and the second to augment it by expanding human abilities exploiting the power of computers in order to augment human decision making. Pragmatists mention chess playing where the present world champion is neither a human nor a computer but rather humans using laptop computers (Baraniuk, 2015). Their view is that we could learn to exploit the power of computers to augment our own skills and stay always a step ahead of AI, or at least not be at a disadvantage. The pragmatists also believe that in the worst of cases a chip can be placed in all thinking

machines/robots to render them inoperative in case of any danger. By concentrating research efforts on intelligence augmentation, they claim we can avoid or minimize the possible danger of AI while providing the means to stay ahead in the race against thinking machines and smart robots.

The doubters: The doubters do not believe that AI is possible and that it will ever become a threat to humanity. Dreyfus (1972), its major proponent, argues that human intelligence and expertise cannot be replicated and captured in formal rules. He believes that AI is a fad promoted by the computer industry. He points out to the many predictions that did not materialize such as those made by Herbert A. Simon in 1958 that “a computer would be the world’s chess champion within ten years” and those made in 1965 that “machines will be capable within twenty years, of doing any work a man can do” (Crevier, 1993). Dreyfus claims that Simon’s optimism was totally unwarranted as they were based on false assumptions that human intelligence is based on an information processing viewpoint as our mind is nothing like a computer. Although, the doubters’ criticisms may have been valid in the last century, they cannot stand for the new developments in AI. Deep Blue became the world’s chess champion in 1997 (missing Simon’s forecast by twenty one years) while we are not far today from machines being capable of doing all the work that humans can do (missing Crevier’s prediction by about a decade. There are also self-driving vehicles, nurse-robots taking care of the elderly, and Google Search which may know better than us what we are looking for. These technologies are hard to explain by those that doubt AI technologies will materialize and affect us profoundly.

A more sophisticated attack comes from doubters who state that it is wrong to believe that once computers have been provided with sufficiently advanced algorithms, they will be able to improve and then replicate the way our mind works. According to them (Jankel, 2015) computers will not be able to achieve the highest of human ability that of being creative as doing so requires breaking the rules and being anti algorithmic. In other words creative breakthroughs cannot be predicted so any algorithm developed by AI to do so will fail, leaving a big vacuum to the sole province of the human mind that is will always stay more valuable than all algorithmic AI technologies put together. It is like the paintings by the masters which are by far superior to those of the millions of average painters put together. This would mean that all tasks requiring creativity, including innovative breakthroughs, strategic thinking, entrepreneurship, risk taking and similar ones would never, or at least not in the foreseeable future, could be done algorithmically, providing humans with a clear superiority versus intelligent machines.

The timing and expected impact of the AI Revolution: Kurzweil predicted that computers will reach human intelligence around 2029 (Kurzweil, 2005) while Singularity will come by 2045. In 2009 Barrat and Goertzel (2011) asked the participants of an Artificial General Intelligence (AGI) Conference to answer the question: “I believe that AGI (however I define it) will be effectively implemented in the following timeframe”. The answers of the 60 participants and the percentages of their answers are shown below:

<u>Time Frame</u>	<u>Response Percent</u>	<u>Response Count</u>
Before 2030	43.3%	26
2030 – 2049	25.0%	15

2050 – 2099	20.0%	12
After 2100	10.0%	6
Never	1.7%	1

With more than two thirds of respondents predicting that AGI will occur before 2050.

The second prediction being asked was: *“I believe that AGI (however I define it) will be a net positive event for humankind”*. Among the 60 participants, 51, that is 85% answer “Yes” while the remaining 15% said “No”.

In a similar survey Muller and Bostrom (2013) asked hundreds of AI experts at a series of conference the following: *“For the purposes of this question, assume that human scientific activity continues without major negative disruption. By what year would you see a (10% / 50% / 90%) probability for such High Level Machine Intelligence (HLMI) to exist?”*

The median answer for the 10% probability was 2022, the 50% probability was 2040 and the 90% probability the year 2075. The timing from the answers of the two surveys is not far apart from those of Kurzweil (although they may have been influenced by them), agreeing that AGI or HLMI are not so far off and with the majority of scientists believing that AI will have a positive effect for humankind.

In a new survey conducted in early March 2016, Etzioni (2016) posed the following question:

“In his book, Nick Bostrom has defined Superintelligence as ‘an intellect that is much smarter than the best human brains in practically every field, including scientific creativity, general wisdom and social skills.’ When do you think we will achieve Superintelligence?”

The answers of the 80 responders (a 41% rate) are summarized below:

0%	In the next 10 years
7.5%	In the next 10-25 years
67.5%	In more than 25 years
25.0%	Never

These dates put the start of superintelligence later than that of Kurzweil or those of the previous surveys but the answers refer to “superintelligence” rather than AI. It is also interesting that only 25% of respondents answer “never”.

Firms and employment: From the Industrial and digital to the AI revolution

The industrial revolution brought far reaching changes to firms and employment while those of the digital continued the decrease of employment in agriculture and manufacturing while contributing to strong increases in services, in particular in computers, the Internet and the mobile phone markets. The digital revolution also resulted in the decrease of the large industrial firm. The expected changes being brought by AI technologies will be just as, or even more significant as those of the Industrial revolution and much harder to predict for two reasons. First, they will depend on the speed that AI technologies will succeed in automating mental tasks currently performed by humans and replacing them in the process, and secondly the extend of the accelerated process of technological change as intelligent computer programs will become available and capable of developing new programs on their own. There is no doubt, therefore, that AI technologies coupled with the exponential

growth of the Internet will affect how firms operate, how they sell their products/services as well as how they are managed, influencing employment patterns.

In 1995 when the digital revolution was in its infancy, in a **Newsweek** article in February of this year, Clifford Stoll wrote *“Baloney. Do our computer pundits lack all common sense? The truth is no online database will replace your daily newspaper, no CD-ROM can take the place of a competent teacher and no computer network will change the way government works”* (Stoll, 1995). Yet, twenty years later most people access newspapers electronically, we have Google search, Amazon is the 18th largest company in the **Fortune** 2016 list while Facebook had more than 1.8 billion users at the end of 2016. The extent of change can be seen in Table 2 that shows the ten largest firms in Fortune 500 list in 1995 and 2016. The transformation is fundamental, demonstrating how the US economy went from the industrial age to a varied structure of firms in retail, oil and gas, hi tech, finance and health with the two largest industrial companies that were number one and two in 1995 to have fallen to the eighth and ninth position in 2016. Moreover, we see Walmart, a traditional retailer, at the top of the list which is interesting for two reasons. First, it confirms the shift from the industrial to a consumer oriented era, and secondly it proves that a traditional, brick-and-mortar firm can adapt to the digital revolution by successfully competing with Amazon (Petro, 2016), one of the most successful digital firms, on both low prices as well as superior customer service.

Table 2: Fortune 500: Largest USA Firms in 1995 and 2016				
1995			2016	
1	GM	Industrial	1	Walmart Retail
2	Ford	Industrial	2	Exxon Oil & Gas
3	Exxon	Oil & Gas	3	Apple Hi Tech
4	Wal-Mart	Retail	4	Berkshire Finance
5	AT&T	Telecom	5	McKesson Health
6	GE	Industrial	6	United Health Health
7	IBM	Computer	7	CVS Health Health
8	Mobil	Oil & Gas	8	GM Industrial
9	Sears Roebuck	Retail	9	Ford Industrial
10	Altria	Industrial	10	AT&T Telecom
Summary			Summary	
4	Industrial		1	Retail
2	Oil & Gas		1	Oil & Gas
2	Retail		1	HI tech
1	Telecom		1	Financial
1	Computer		3	Health
			2	Industrial
			1	Telecom

What is even more interesting than just “largeness” is the market capitalization of digital versus traditional firms. Table 3 lists four digital firms and four traditional ones comparing their market capitalization, revenues and number of employees. There are some important differences between the two groups. The number of employees of digital firms is about 15% of those of traditional ones while the market cap of the former is more than 67% that of the latter, whereas the revenues of traditional firms is more than 218% that of digital ones. Given the significantly lower number of employees, the higher market cap and the lower

revenues of digital firms, the market cap per employee of high tech firms is an astonishing \$4.13 million while their revenues per employee is \$0.95 million versus \$0.38 and \$0.31 million respectively for traditional firms. There are probably several reasons for such large differences, three major ones being great productivity improvements through the extensive use of digital technologies including the Internet, continuous innovation of the products/services being offered to consumers and of course, high expected growth rates in their future revenues, skyrocketing their share price. For instance, Facebook's market cap per employee of \$22.1 million is more than nine times that of J&J's \$2.43 million. But even its revenues per employee of \$1.15 million are more than ten times higher than that of J&J, demonstrating the elevated productivity of Facebook's workforce.

Table 3: Market Capitalization and Revenues: Digital and Traditional Firms*						
	Company	Number of Employees (000)	Market Cap (Billion)	Market Cap per Employee (Million)	Revenues (Billion)	Revenues per Employee (Million)
Market Cap: Digital Firms	Apple	115	617.5	5.37	233.7	2.03
	Google	57	560.8	9.84	74.5	1.31
	Amazon	270	365.9	1.36	107	0.40
	Facebook	16	347.2	22.1	18	1.15
	Total	458	1,891	4.13	433.2	0.95
Market Cap: Traditional Firms	Berkshire	331	407.2	1.23	210.8	0.64
	J&J	127	308.9	2.43	15.4	0.12
	Walmart	2,200	214.6	0.10	482	0.22
	Toyota	346	200.7	0.58	237	0.68
	Total	3,004	1,131	0.38	945	0.31
*The information for these firms was collected on 12/12/2016						

The four digital firms are at the forefront of AI, investing huge amounts of money for internal AI research as well as buying promising start ups (CB Insights, 2016; Metz, 2016) with the purpose of offering AI technologies to their customers as well as developing new applications for the market place. Their huge market capitalization is a definite factor allowing them to acquire whatever start up is considered valuable for them to stay ahead in the AI race. The interesting question is if these four digital firms will remain at the top twenty years from now when the AI revolution will be the dominant force determining success and determining the high value added products and services.

The successful, dominant firm of the AI revolution and its management

Revolutions, by definition, are associated with major changes. The Industrial one brought the large, industrial firm that exploited the power of machines to substitute, supplement and amplify the manual work performed by humans, increasing productivity considerable and offering affordable products to consumers, significantly increasing market size and living standards. The digital revolution exploited the power of computers to substitute, supplement and amplify the routine mental tasks performed by humans also improving productivity and further contributing to reduced prices. As previously mentioned, the AI revolution aims to substitute, supplement and amplify **practically all** tasks currently performed by humans, becoming in effect, for the first time, a serious competitor to them.

Accepting this watershed reality, setting aside criticisms about AI similar to those levied in 1995 against the Internet (Stoll, 1995), and accepting our short-sightedness to predict bold technological inventions, I will attempt, nevertheless, to predict the dominant firm twenty years from now and the unavoidable shift in employment.

In 1995 from the four digital firms listed in Table 3 two (Google and Facebook) were not even founded, Amazon was just a year old while Apple was nineteen years old but in a grave financial state, with a profit of just \$0.4 billion that year and a loss of \$0.8 and \$1 billion in 1996 and 1997. It was not possible, therefore, to have predicted in 1995 that the market capitalization of these four firms would be (at the end of 2016) close to \$1.9 trillion (more than the GDP of the 120 poorest countries in the world) or that their market cap per employee would be \$4.13 million. Similarly, there is no accurate way to currently predict the dominant, successful firms of 2037 or the critical factors leading to such success as it was impossible to do so in 1995. Needless to say, we can consider these and extrapolate them assuming that they may still hold true between now and 2037.

First and foremost all four firms have been extremely innovative, each one in its own unique way. Second, they have all used the Internet in a super effective manner to provide their services, sell their products and streamline their operations. Third, they have been successful in hiring top talent and motivate their top employees with a pleasant work environment, high salaries and generous stock options. Finally, these four firms have grown significantly by acquiring other companies, often promising start-ups in the selective areas they want to expand or attain expertise. Interestingly none of these four characteristics can be automated, or become part of an algorithm, at least during the next twenty years. In my view, they will continue to remain critical factors for succeeding in the future and will depend greatly on people's decisions and actions to implement them. Furthermore, algorithms will have to be modified as competitive, market, environment and other factors will be changing and only humans will be able to identify when such changes have occurred (otherwise, the algorithms will be dysfunctional). Below each of these four success factors is discussed in further detail.

Innovative breakthroughs: There are different types of innovation from minor, marginal ones to distinctive breakthroughs as those implemented by the four digital firms of Table 3. There are numerous books and articles written about innovative breakthroughs like the iPhone, or Google search and how they have transformed the world and brought riches to those that invented them. But what must be understood is that such breakthroughs are exceptional and only identified as such after the fact. In an older but fascinating book about breakthroughs (Ketteringham and Nayak, 1986), its authors concluded that there was one common characteristic of all the sixteen cases they studied. Nobody believed that they would be successful, urging their developers to give up wasting their time and company resources. These breakthroughs succeeded only because of the persistence of their inventors that did not give up despite all hardship and continued their development often in their spare time. The value of breakthroughs is hard to appreciate as they involve brand new applications people are not familiar with, the value of which they cannot understand, or do not believe as technically possible as was the case with fibre optics. Even the value of Google search was not apparent in the late 1990s as can be attested by the fact that its inventors could not sell it for the asking price of 1.6 billion (Battelle, 2005).

It cannot be assured that by identifying interesting innovative ideas they will turn into commercially successful applications. The road is long and full of risks as many things can go wrong, including technological problems, inability to ensure adequate financing, competitors developing similar ideas and possible delays until people realize their importance and decide to use them. Moreover, it is worth mentioning that the great majority of innovative ideas do not succeed in a big way, that new competitors can improve the original idea and capture market shares and that many inventions can become obsolete by more novel ones often introduced by garage start-ups (see below). Creative destruction is at the heart of market economies and the major contributor to productivity increases as well as the offering of a greater choice of product/services to consumers. No concrete advice on how to innovate, not to mention achieve breakthroughs, can be provided beyond saying that some organizations are better than others, at least for some period of time, in creating the right environment to cultivate innovation and exploit its benefits. It must be emphasized, therefore, that success is unpredictable and that the more valuable the innovative idea the harder it is to conceive and implement.

Technologies and their Usage: The digital revolution has provided significant productivity improvements for the back office operations of firms and allowed consumers to buy goods and services online without having to physically go to a store or office. In the process it has eliminated a large number of repetitive mental tasks performed by people and shifted employment patterns. This is evident in Table 3 by comparing the number of employees between digital and traditional firms. For instance, the number of employees of Walmart is more than eight times higher than that of the other retailer Amazon, while the revenues of the former is only 4.5 higher than the latter, showing the great productivity improvements being realized by Amazon's effective utilization of digital technologies, and the Internet. It is likely that the AI revolution will continue or even accelerate that pattern, further improving productivity and reducing employment with some firms being more successful than others by exploiting emerging technologies and investing considerable sums to do so.

The uniqueness of the AI technologies is their potential to supplement, substitute and amplify **practically all** tasks currently performed by humans with critical consequences for firms that must achieve significant productivity improvements to stay competitive, but at the same time raising the possibility of increased unemployment (see below). Extrapolating from the impact of the Industrial and digital revolutions it seems that technology has created more jobs than it has destroyed (Stewart, Debapratim and Cole, 2015) although there may be a transitional period of increased unemployment until new opportunities are created to serve the emerging needs of those with increased incomes. The difficulty would be in knowing which of the AI technologies would provide the greatest benefits and determining the right time to start investing in them. It is doubtful that such an investment decision could be made algorithmically as it would require predictions about the pros and cons of future and uncertain technologies. This means that top executives will have to consider the advantages and drawbacks of the available investment possibilities and make the final decision(s) judgmentally using all information at hand.

Managing people: It is also highly likely that the trend towards a smaller payroll will continue as AI technologies will accelerate the number of tasks that can be performed by

machines and robots. The more jobs being automated however, the greater the skills that would be required to adequately perform the remaining tasks, for both the efficient operation of firms as well as for utilizing AI and other technologies in the best possible way. This would require hiring talented employees and motivating them to get the most out of their performance in order to attain and maintain competitive advantages over other firms. The major differentiator would therefore come from these talented individuals conceiving and implementing innovative ideas and winning strategies that would steer the organization onto a successful path. Hiring, motivating and successfully managing talented individuals would probably be one of the most critical success factors for firms in the AI era and would also be impossible to put into an algorithm.

Growth by acquisition: The four digital firms of Table 3 have grown substantially through acquisitions which also allowed them to obtain and secure desired expertise in vital areas. Together they have acquired more than 400 firms (Whittick, 2016) with Google leading the way with 190, Apple with 82, Amazon with 67 and Facebook with 62 (including WhatsApp bought for \$19 billion, a huge price for a start up with only 55 employees and less than five years old). In the past, new technological inventions have fundamentally affected the revenues and profitability of firms and have become a critical success factor. It is exceedingly likely that this pattern will continue, if not accelerate with AI technologies creating uncertainty for established firms as such technologies, or who will introduce them first cannot be predicted. The huge market capitalization of the four firms of Table 3 makes it easier for them to acquire start-ups which they believe will be strategically important for their future expansion through the issuing of more shares if necessary. In the process they can maintain and/or increase their monopolistic/oligopolistic advantages while also squashing possible, future competition. It is interesting to consider the implications if this trend continues in the future and what it will mean for competition as it can create a “winner takes all” environment favouring today’s dominant firms.

Competitive Requirements: Big Data, Algorithmic Decisions and Operational Excellence

Let us assume that an advanced algorithm is developed to predict exactly what customers want. What will the result be? Even if it is proprietary, competitors will inevitably figure out how it works and also utilize its predictions after some period of time. If it is available to everyone at a price, all competitors will swiftly adopt this algorithm to stay competitive. Those who do not utilize it will probably go out of business given its superiority to predict customers’ wants. There will be no special benefit as the advanced algorithm would simply become a competitive requirement for staying in the race. Big data (Schonberger and Murray, 2014) offers the opportunity of developing successful algorithms to understand what customers want and as such can be extremely useful for decision makers. The challenge is that both the data and the techniques to analyse them are available to practically everyone interested, turning their recommendations into a competitive requirement. Jankel (2015) is therefore probably right, stating that “computers will never create disruptive innovations” as well as be able to provide advice on the other three major factors (selecting the right technologies, managing people and making the appropriate acquisitions) characterizing the four digital firms of Table 3.

The 2037 Successful Firm

The successful 2037 AI era firms cannot be predicted but they will probably be closer to the digital than the traditional ones of Table 3. They will exploit global opportunities to the limit in order to design, produce and sell their products/services and they must be willing to assume entrepreneurial risks to innovate and grow. The speed of technological change from the forthcoming AI revolution will open huge opportunities for growth and profitability but also new challenges and competition from new garage type start-ups as breakthrough ideas can come from anywhere and their development and financing will be easier through crowd sourcing and venture capital. Only time will tell if the successful firms of the digital revolution (as the four listed in Table 3) will still be dominant in 2037 or if they will be replaced by new ones. In my opinion, there are accrued competitive advantages that could be exploited to stay at the top, but at the same time big technological breakthroughs can fundamentally change the competitive landscape and wipe out such advantages. But what is probable is that AI will bring revolutionary changes to the business environment and the way firms are inspired, founded and managed at a much greater scale than the digital era. The successful firms during the AI revolution will have in addition to the Chief Innovation Officer (CIO) to also appoint a Chief Artificial Intelligence Officer (CAIO) to be in charge of evaluating and exploiting AI technologies to gain the most out of their implementation in all aspects of the firm.

Garage start-ups, crowd sourcing and VC Funding



The way the four digital firms of Table 3 were created was quite different than that of traditional ones. They were conceived with a vision to change the world, even though making money was no doubt also part of their founders' objective. Financing was secured mainly through Venture Capital (VC) (with the exception of Apple that was founded well before VC funding was readily available) that invested heavily in these firms that managed to grow rapidly over time (more than 20% a year) and achieve stellar market capitalization putting all four of them at the top ten in the USA. This is quite different to that of traditional firms that depend heavily on the personal financing of their owners, supplemented with bank loans, and apart from a few exceptions, experience low growth rates, taking decades to reach the top.

The widespread utilization of the Internet has provided two significant tools for start-ups that can utilize crowd funding and crowd creativity to improve their chances of success. Such tools coupled with VC funding will increase the number of new firms aimed at exploiting AI technologies and the chance to succeed in big ways like Google, Amazon or Facebook. Breakthrough ideas can come from anywhere and they do not require expensive laboratories or huge financing to develop and market. This means that garage operations, small offices in incubators or accelerators, as well as cubicles in universities, can provide an initial base for start-ups that can emerge everywhere outside Silicon Valley as their value in wealth creation is becoming obvious. Today a large number of countries and individuals understand the importance of innovation and entrepreneurship for their economies and encourage their youth to implement their innovative ideas by starting their own business. Israel has been extremely successful in doing so (Senor and Singer, 2009) as has China, Germany, UK, France, Chile and New Zealand among others. Characteristic of efforts to encourage and expand start-ups is France's latest mega-35,000 m² campus being built in Paris to house start-ups under one roof, or the vision of Chinese's President Xi Jinping to turn his country into a global technology powerhouse. The implications of these efforts will

be the globalization and acceleration of technological innovation, greater competition and more equal distribution of employment and wealth across nations that encourage technological innovation.

Employment patterns

The 1995 paper showed the continuous decline in the agricultural and manufacturing (since it peaked in the 1950s) employment and the growth in that of services. Table 4 and 5 indicate that the same pattern has continued since 1995 with employment dropping from 2.5% to 0.7% in agriculture and from 17% to 14% in manufacturing in the USA and correspondingly from 2.1% to 1.3% and from 21% to 15% in the UK while witnessing increases in the percentage of employment in services. If current trends continue agriculture and manufacturing employment will further decrease as AI automation will affect both agricultural and manufacturing operations and jobs. This means that employment in services will have to broaden to compensate for the jobs being lost in agriculture and manufacturing. However, the service sector will also witness considerable realignments as many jobs will be eliminated and new ones will be created.



What critics fear is the speed of job obsolescence through AI technologies in particular in the service sector. They say that while it took more than two centuries to witness the full impact of the Industrial revolution and three or four decades to experience that of the digital one, it may be no more than a decade until we observe the full effects of the AI revolution (The New Yorker, 2016). Another concern is the extent of displaced jobs that some studies (Frey and Osborne, 2013) estimate as 47% of all occupations. What experts are afraid of is whether the social structure could withstand such a fast and huge reduction in jobs. They claim that even if there were vast retraining programs it may not be easy to avoid societal disruptions as the new jobs would require skills that may not be so easily attainable. In addition to the truck and taxi drivers already mentioned that could be replaced by self-driving vehicles, there are many other jobs at stake. According to PBS (Thoet, 2016) the newly announced Amazon Go retail store, using AI technologies to abolish employees, *“could drastically change the way people shop and eventually eliminate the need for millions of workers, industry experts predict.”* Other jobs that could be affected range from preparing and taking payments for fast food to skilled professions including financial advisers, medical specialists as well as high level IT tasks that could be automated using algorithms such as those available through Amazon’s Web Services. The big question is where will all these lead and what will the implications be for employment and wealth distribution?”

A recent McKinsey study entitled *“Where machines could replace humans – and where they can’t (yet)”* (Chui, Manyika and Miremadi, 2016) distinguishes three groups of occupational activities that are highly susceptible, less susceptible and least susceptible to machines and robots taking over jobs currently performed by humans (with the number in parentheses denoting the estimated percentage substitution). The first group includes “managing others” (9%) and “applying expertise” (18%), the second covers “interactions among stakeholders” (20%) and “unpredictable physical work” (25%) while the third contains “data collection” (64%), “data processing” (69%) and “predictable physical work” (78%). The above classification is consistent with other studies which found that jobs related to social skills grew 10% a year between 1980 and 2012 while all others declined by 3% during the same period (Deming, 2016). However, will the new jobs being created compensate for, or even augment those lost by technology as in the past? Opinions differ greatly on this issue as some are arguing that there is a fixed supply of work to be done, while others, with Milton Friedman the most notable advocate, claims that human wants and needs are infinite and that clever entrepreneurs will always devise ways to fulfil them as long as sufficient buying power is available. Others worry, however, saying “what if the jobs of the future are also potentially automatable?” (The New Yorker, 2016).

In an article entitled “Will Humans Go the Way of Horses? Labour in the Second Machine Age” (Brynjolfsson and McAfee, 2015) its authors consider both sides of the argument and conclude that there are no obvious answers. Instead they propose:

“It’s time to start discussing what kind of society we should construct around a labor-light economy. How should the abundance of such an economy be shared? How can the tendency of modern capitalism to produce high levels of inequality be muted while preserving its ability to allocate resources efficiently and reward initiative and effort?”

What do fulfilling lives and healthy communities look like when they no longer center on industrial-era conceptions of work?"

The critics claim that even if technological automation may not increase unemployment, it can destroy middle range jobs while increasing those on the low and high ends, augmenting social inequality as the pay between low and high end jobs is amplified. Governments must, therefore, enact policies to minimize such inequalities either by guaranteeing a minimum universal income, raising taxes for the superrich and/or increasing inheritance taxes. During the digital revolution the demand for computer and data specialists rose steeply their salaries while eliminating large numbers of clerical jobs. The prediction is that the same thing is happening for AI specialties, such as deep learning professionals, necessary for the further development of AI technologies (Metz, 2016). In all likelihood such demand will further intensify until more people are educated to be able to fill these jobs.

In addition to specialized AI professions, employment opportunities will exist by moving from traditional jobs susceptible to automation, to those demanding social and interpersonal skills as well as creativity and innovation. There will also be a demand for novel jobs that will aim to satisfy the needs of higher income people that can afford personal trainers, coaches, beauty advisors, diet consultants, nutritionists and teachers for their children, among others. These and similar jobs will offer employment opportunities and adequate payment, and will compensate for those lost when AI technologies will replace existing jobs.

Finally, the impact of the AI revolution will probably be more pronounced in developing countries than in advanced ones for two reasons. First, as unskilled and semiskilled labour will be replaced by computers and robots there will be no reason for firms to move their production to developing nations to exploit their cheap supply of labour as they can achieve the same or cheaper costs utilizing AI technologies, thus increasing the trend towards "reshoring" back to advanced countries (Ford, 2016). Second, developing countries will be at a disadvantage by not being able to invest in expensive AI technologies, particularly since such technologies will reduce the demand for human labour thus further increasing unemployment. Unfortunately, there will be no easy solutions, with the greatest challenge being how to educate their young people in AI technologies and by doing so become able to attract investments from abroad. Otherwise, they will stay behind as the nations that did not manage to industrialize their economies.

Work and Leisure

Machines, electrification, cars, computers, the Internet and smartphones among others inventions have affected our lives and shaped our work and social environment. The impact of AI technologies can be even more profound than that of both the Industrial and digital revolutions put together, as it holds the potential to affect practically all tasks currently performed by humans, diminishing the amount of work left for people and increasing wealth inequalities and their free time. Proponents of the AI technologies see this as a positive development liberating people from routine work, allowing them to pursue their own interests. The critics say that this will further increase inequalities as fewer people will hold the well-paying jobs and the majority will depend on part time work or limited employment opportunities and therefore receive a lower income. There are no obvious answers to this argument and it will all depend as Brynjolfsson and McAfee say "*what kind*

of society we should construct around a labor-light economy? How should the abundance of such an economy be shared?" At the beginning of the Industrial revolution people used to work fifteen hours a day, five days a week. Today a standard work week is less than half and there is no reason that it cannot be halved in the future as productivity increases.

Aristocrats did not work at all in the past, devoting all their time to leisurely activities, hobbies, holidays and travel and things they were interested in. The citizens of ancient Athens spent practically all their time philosophising, exercising and concerned with democracy, while slaves did all the work. The optimists dream is that we can all become the "new" aristocrats, or the "modern" Athenians, with computers and robots as our slaves doing all the housework, the shopping and working at the office, the factory or land. Whether this dream is a utopian or dystopian future is left up to the reader to decide, not underestimating however, that intelligent machines will eventually become at least as smart as us and a serious competitor to the human race if left unchecked and if their great potential to augment our own abilities is not exploited to the maximum.

Conclusions

The societal impact of the digital revolution has been significant as it has affected most aspects of our lives and work, having molded the dominant firm, shaped our shopping and entertainment habits as well as our employment patterns. This paper argues that the AI revolution is on target and will come into full force within the next twenty years as did the digital one since 1995 and will probably have an even greater impact than both the Industrial and digital ones combined. What is uncertain is if such an impact will lead to a utopian or dystopian future, or somewhere in between, which according to the optimists will happen when some mutually beneficial coexistence is achieved by accepting our limitations and seeking ways to augment our own decision making abilities in a similar way that the world chess champion is now a human utilizing the help of a computer to determine his moves. The pessimists worry that when AI machines become smarter than us they will be making all our decisions for us, from the less important ones like SIRI or Cortana choosing a restaurant and ordering food and wine for their owners, to important ones like driving cars or managing nuclear plants. Pessimists are also concerned with social discontent as the amount of work available for people will diminish, and with the increasing wealth inequality between those employed or owning AI applications and the rest.

This paper has described the substantial uncertainty about the future impact of AI technologies and their potential to create a utopian or dystopian world. As we head towards uncharted waters and uncertain choices the obvious challenge is what can be done to maximize the chances of exploiting the benefits while avoiding the negative consequences of AI technologies? There are two positive aspects in dealing with this challenge. First, the dangers are well understood and secondly, there is plenty of time to debate the issues and take wise actions to deal with them effectively. Increased unemployment and greater wealth inequality are debatable given that the Industrial and digital revolutions decreased rather than increased unemployment (Stewart, Debapratim and Cole, 2015), although some claim that this may change drastically with the widespread introduction of AI technologies leading to massive job reductions and bringing us towards Huxley's ***Brave New World***. Scientists like Etzioni (2016), however, do not believe that AI is a threat to humanity, maintaining that it would be a great pity to forfeit its great advantages out of

unsubstantiated fears that it may get out of control. There is little doubt that new technological breakthroughs similar to those of the Internet and smartphones of the digital revolution will probably emerge during the next twenty years and that will greatly impact our societies, lives and work in general. The challenge will be to identify such breakthroughs as early as possible in order to exploit their benefits.

Concerns about the AI related risks are not unique. Possible catastrophes, in particular those that threaten to destroy the environment or end human civilization attract people's attention and are popular among both scientists seeking to publicize their work and journalists looking for "cool" stories. Nuclear wars, global warming, accidents in nuclear plants, unchecked epidemics, and geoengineering, among others, worry people who are concerned about their implications. At the same time it is clear that AI risks cannot be ignored even though the probability of their occurrence may be extremely small as their potential impact, according to critics, can be devastating, ending human supremacy when machines smarter than people are developed. At the same time progress cannot be halted which means that the only rational alternative is to identify the risks involved and devise effective actions to avoid their negative consequences. In ending this paper I would like to quote Faggella (2016) saying: *"Thinking about the risks associated with emerging AI technology is hard work, engineering potential solutions and safeguards is harder work, and collaborating globally on implementation and monitoring of initiatives is the hardest work of all. But considering all that's at stake, I would place all my bets on the table and argue that the effort is worth the risk many times over"*.

REFERENCES

- Baraniuk, C. (2015). The Cyborg Chess Players that Cannot Be Beaten, *BBC News*
<http://www.bbc.com/future/story/20151201-the-cyborg-chess-players-that-cant-be-beaten>
- Barrat, J., (2013). *Our Final Invention: Artificial Intelligence and the End of the Human Era*, Thomas Dunne Books
- Barrat, J. and Goertzel, B. (2011). How Long Till AGI? – Views of AGI-11 Conference Participants, *HPlusMagazine.com*
<http://hplusmagazine.com/2011/09/16/how-long-till-agi-views-of-agi-11-conference-participants/>
- Battelle, J. (2005). *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*, New York, NY: Penguin
- Best, J. (2016). IBM Watson: The inside story of how the Jeopardy-winning supercomputer was born, and what it wants to do next, *TechRepublic.Com*
<http://www.techrepublic.com/article/ibm-watson-the-inside-story-of-how-the-jeopardy-winning-supercomputer-was-born-and-what-it-wants-to-do-next/>
- Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*, Oxford University Press, Oxford
- Brynjolfsson, E. and McAfee, A. (2015). Will Humans Go the Way of Horses?, *Foreign Affairs*
<https://www.foreignaffairs.com/articles/2015-06-16/will-humans-go-way-horses>
- CBInsights.Com (2016). The Race for AI: Google, Twitter, Intel, Apple in A Rush to Grab Artificial Intelligence Startups
<https://www.cbinsights.com/blog/top-acquirers-ai-startups-ma-timeline/>
- Chouard, T. (2016). The Go Files: AI computer clinches victory against Go champion, *Nature.Com*
<http://www.nature.com/news/the-go-files-ai-computer-clinches-victory-against-go-champion-1.19553>
- Chui, M., et al. (2016). Where machines could replace humans – and where they can't (yet), *McKinsey.Com*
<http://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/where-machines-could-replace-humans-and-where-they-cant-yet>
- Clark, J. (2015). Why 2015 Was a Breakthrough Year in Artificial Intelligence, *Bloomberg.Com*
<https://www.bloomberg.com/news/articles/2015-12-08/why-2015-was-a-breakthrough-year-in-artificial-intelligence>

Crevier, D. (1993). *AI: The Tumultuous History of the Search for Artificial Intelligence*, Basic Books

Deming, D.J. (2016). *The Growing Importance of Social Skills in the Labor Market*, NBER working paper series

Dreyfus, H. L. (1972). *What Computers still can't do: A Critique of Artificial Reason*, MIT Press

Etzioni, O. (2016). No, the Experts Don't Think Superintelligent AI is a Threat to Humanity, *MIT Technology Review*

<https://www.technologyreview.com/s/602410/no-the-experts-dont-think-superintelligent-ai-is-a-threat-to-humanity/>

Faggella, D. (2016). Exploring the risks of artificial intelligence, Mar 21, *Crunch Network NETWORK*

<https://techcrunch.com/2016/03/21/exploring-the-risks-of-artificial-intelligence/>

Ford, M. (2016). *Rise of Robots: Technology and the Threat of a Jobless Future*, Basic Books, New York

Frey, C. B. and Osborne, M. A. (2013). *The Future of Employment: How Susceptible are Jobs to Computerisation?*, Oxford Martin School

http://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf

Harari, Y. N. (2016). *Homo Deus: A Brief History of Tomorrow*, Harvill Secker, New York

Cellan-Jones, R. (2014) Stephen Hawking warns artificial intelligence could end humanity *BBC News*

<http://www.bbc.com/news/technology-30290540>

Jankel, N. S. (2015). AI vs. Human Intelligence: Why Computers Will Never Create Disruptive Innovations, *Huffington Post*

http://www.huffingtonpost.com/nick-seneca-jankel/ai-vs-human-intelligence-b_6741814.html

Joy, B. (2000). Why The Future Doesn't Need Us, *Wired.Com*

<https://www.wired.com/2000/04/joy-2/>

Kahn, J. (2016). Google's DeepMind AI Takes on Popular Video Game Starcraft, *Bloomberg.Com*

<https://www.bloomberg.com/news/articles/2016-11-04/deepmind-master-of-go-takes-on-popular-video-game-starcraft>

Kettingham, J. M. and Nayak, P. R. (1986). *How the Vision and Drive of Innovators in Sixteen Companies Created Commercial Breakthroughs that Swept the World*, Rawson Association

Kurzweil, R. (2005). *The Singularity is Near: When Humans Transcend Biology*, Viking Press

Makridakis, S. (1995). The Forthcoming Information Revolution: Its Impact on society and firms, *Futures*, Vol. 27, No. 8, pp. 799 – 821

Markoff, J. (2016). *Machines of Loving Grace: The Quest for Common Grounds Between Humans and Robots*, Harper Collins, New York

May, K. T. (2013). How did supercomputer Watson beat Jeopardy champion Ken Jennings? Experts discuss, *TedBlog* <http://blog.ted.com/how-did-supercomputer-watson-beat-jeopardy-champion-ken-jennings-experts-discuss/>

Metz, C. (2016). Google, Facebook and Microsoft are remaking themselves around AI, *Wired.Com* <https://www.wired.com/2016/11/google-facebook-microsoft-remaking-around-ai/>

Muller, V. C. and Bostrom, N. (2013). *Future Progress in Artificial Intelligence: A Survey of Expert Opinion*, Fundamental Issues of Artificial Intelligence, Synthese Library, Berlin: Springer

New Yorker (2016). Our Automated Future: How long will it be before you lose your job to a robot? December 19 & 26, 2016

Parloff, R. (2016). Why Deep Learning is Suddenly Changing Your Life, *Fortune* <http://fortune.com/ai-artificial-intelligence-deep-machine-learning/>

Peckham, M. (2016). What 7 of the World's Smartest People Think About Artificial Intelligence, *Time.Com* <http://time.com/4278790/smart-people-ai/>

Petro, G. (2016). Amazon Vs. Walmart: Clash of the Titans. *Forbes.Com* <http://www.forbes.com/sites/gregpetro/2016/08/25/amazon-vs-walmart-clash-of-the-titans/#2f18402d18d9>

Say, J. B. (1828). *Cours complet d'économie politique*, Chez Rapilly, Paris.

Senor, D. and Singer, S. (2009). *Start-up Nation*, Hachette Book Group

Schonberger, V.M. and Cukier, K. (2014). *Big Data: A Revolution That Will Transform How We Live, Work, and Think*, John Murray, London

Stewart, I., Debapratim, D. and Cole, A. (2015). *Deloitte* <https://www2.deloitte.com/content/dam/Deloitte/uk/Documents/finance/deloitte-uk-technology-and-people.pdf>

Stoll, C. (1995). " Why the Web won't Be Nirvana", *Newsweek* <http://europe.newsweek.com/clifford-stoll-why-web-wont-be-nirvana-185306?rm=eu>

TheWeek.Com, (2016). Google has more than 1,000 artificial intelligence projects in the works. <http://theweek.com/speedreads/654463/google-more-than-1000-artificial-intelligence-projects-works>

Thoet, A. (2016). What Amazon's cashier-free store could mean for millions of workers, *PBS.Org*
<http://www.pbs.org/newshour/rundown/amazons-cashier-free-store-mean-millions-workers/>

Whittick, S. (2016). We visualized acquisitions by the Big Five tech companies since 1985, here's what we learned, *Geckoboard*
<https://www.geckoboard.com/blog/acquisitions-mergers-big-five-tech-companies/#.WFkFPVN96M9>