Artificial Intelligence: Uses and Misuses

By Hibatullah Alzahrani

Saudi Arabian Cultural Mission

Abstract- Artificial Intelligence (AI) was mostly regarded as science-fiction in the past but with the recent advancements in technology, it has silently crept into our lives. From social media to computer games to self-driving cars to military gadgets to personal digital assistants, AI is everywhere. This progress is also due to a paradigm shift in AI community where current trend is to make AI stronger in specific domains rather than making a human-like AI which can do anything. Resultantly, AI can now out-perform humans in many areas. But this progress of AI is scary for some people who are predicting the "rise of machines" in half a century or so if AI progress remains unbridled.

Keywords: artificial intelligence, AI, social media, strategy-games, robots, ethical issues, AI future

GCST-D Classification : F.1.1

ARTIFICIALINTELLIGENCEUSESANDMISUSES

Strictly as per the compliance and regulations of:

Artificial Intelligence : Uses and Misuses

Hibatullah Alzahrani

Abstract- Artificial Intelligence (AI) was mostly regarded as science-fiction in the past but with the recent advancements in technology, it has silently crept into our lives. From social media to computer games to self-driving cars to military gadgets to personal digital assistants, AI is everywhere. This progress is also due to a paradigm shift in AI community where current trend is to make AI stronger in specific domains rather than making a human-like AI which can do anything. Resultantly, AI can now out-perform humans in many areas. But this progress of AI is scary for some people who are predicting the "rise of machines" in half a century or so if AI progress remains unbridled.

Keywords: artificial intelligence, AI, social media, strategy-games, robots, ethical issues, AI future.

I. Introduction

Artificial Intelligence, referred more commonly by its acronym AI, is one of the most fascinating and most mysterious of the modern technologies. Whenever someone mentions AI, an idea pops up in our mind about a super-intelligent computer – one which can understand what people are saying and respond to them and think autonomously and can obey any command issued by its human master. For some people, the picture is more grotesque as the concept of AI bears into their mind the idea of "terminator" like killing robots, hell bound on eliminating the human race and taking over the world. In reality, AI is still at very nascent stage and no way near replicating the intricate behavior of human brain. But though AI can't outperform humans in general way, it can excel the humans in some specific areas. In 1997, AI shocked the world by defeating the then reigning champion of chess Gary Kasparov and more recently, another AI machine, built using Google DeepMind, has defeated the human champion in much more complex game of Go. But these are specifically built machines, designed to do one and only one task. And that's the way Artificial Intelligence has evolved – trying to be the master in few trades rather than the jack of all. From personal digital assistants like Siri and Cortana to self-driving cars to conversational bots, AI is becoming more and more adept at the tasks assigned to it. Behind its rise is the exponential rise in computational power and storage capabilities which in turn have given rise to complex machine learning models like deep neural nets which are really the power-source behind AI functionality and its learning ability.

It is hard to define what Artificial Intelligence really means – in literal terms we can say that Artificial Intelligence means the *machines with the ability to think*.

Author: Saudi Arabian Cultural Mission. e-mail: mesfer66@gmail.com

But then it becomes important to define what thinking really means and we can get into all kinds of philosophical debates. Alan Turing, pioneer of the Artificial Intelligence, described AI as *the machines which can learn from experience* [1]. But that is only one part of the story, the other part is about taking some definite action as response to external stimuli. Thus, the modern definition of Artificial Intelligence is given as [2]: "*A rational and flexible agent that senses its environment and takes some action which maximizes its chance of success at some predetermined goal*".

II. A Brief History of AI

The field of AI has seen quite a few ebb and flow during its evolution. The recent rise in its popularity can be partially attributed to increase in computational power of the machines and partially to the realization of its limitations by researchers who now set more realistic and achievable goals. It is difficult to estimate who first came up with the idea of intelligent machines but the concept is widely attributed to famous paper of Alan Turing in 1950, titled "Computing Machinery and Intelligence" [1], in which he raised the question *"Can Machines think?"*. He went on to define thinking of machines as *learning from experience* and also proposed the famous Turing Test (which he originally called as Imitation Game) in which a human investigator is asked to converse with a human and a machine remotely and then tell which one is machine. The test has been criticized for being too narrow as it focuses on only one aspect of AI – namely Natural Language Processing, whereas in reality an AI agent can do many more tasks without doing language processing. The field of AI formally came into being in the 1956 Dartmouth Conference [3], organized by Marvin Minsky, John McCarthy, Claude Shannon and Nathan Rochester. John McCarthy was the one to propose and convince others to accept the name "Artificial Intelligence" for this field. The next two decades were considered the golden years of AI, with huge funding coming from various sources, most notably ARPA (then called DARPA, Department of Defense Advanced Research Projects Agency). The people working in the AI were highly optimistic of its success, especially after some early breakthroughs, and one of the pioneers went as far as to say that within twenty years, machines could do anything that a human can do [4]. Their optimism was based on development of fairly complex AI programs like Daniel Bobrow's STUDENT [5], which could solve simple high-school algebraic problems and Herbert Gelernter's Geometry Theorem Prover [6]. John

McCarthy wrote and then refined the Lisp [7] programming language in late 1950s which even today remains the lingua franca of AI world. But these researchers failed to acknowledge the complexity of remaining problems in AI as well as the limitations of computational power at the time which couldn't handle very large and complex problems. In early 70s, AI came under huge criticism for failing to deliver and funding dried up. At the same time, Perceptrons, which were thought of as analogous to neurons in brain, didn't live up to their potential and the idea to replicate the working of human brain failed miserably. Thus started the first AI winter which continued until the early 80s when the "expert systems" began to be widely used in corporations. Expert systems were custom designed for specific domains and used the knowledge of that domain to answer certain questions and solve complex problems. Funding revived for several projects, most notably the Japan's Fifth Generation Computer Project [7], worth almost $1B in 1981, which set-off a chain of similar projects from USA and UK. Around the same time, Perceptrons were back in popularity due to "back-propagation" [8] algorithm. The AI industry was worth billions when it collapsed again in 1987 with the fall of "Lisp Machines". The desktop machines from Apple and IBM became cheaper than specialized AI hardware and the AI market no longer remained lucrative. The second AI winter continued until the biggest event in AI history came – on 11 May, 1997, Deep Blue [9] became the first computer to beat the then reigning chess-champion Garry Kasparov. This single event restored the people's belief in AI; finally AI was able to outmatch the humans in some field. It also taught the AI people an important lesson that focusing on a single problem domain at a time is way more beneficial than trying to build an all-round AI machine which can do anything. Thus sprang a series of astonishing events which strengthened the AI's macho in the tech arena – a Stanford robot won the DARPA Grand challenge by driving autonomously for 131 miles in 2005 [10]; IBM Watson defeated two of the best Jeopardy players in Jeopardy Quiz Show in 2011 [11]; self-driving cars began to perform at near-human levels; Microsoft Kinect was able to capture the gestures of players and gave them a taste of real-world scenario; Apple's Siri and several other chatbots became common which recognized human speech much more accurately and gave responses accordingly; and finally in March 2016, Google's AlphaGo [12] won 4 out of 5 games of Go to defeat the Go champion Lee Sedol. Much of this success is attributed to two things – firstly computers have become way more powerful which have enabled the statistical models to be built using huge amounts of data and fancy algorithms like Deep Learning. Secondly, there has been a change in AI paradigm such that AI is no longer thought of as a self-acting autonomous machine that can do anything which humans can. Rather AI is now defined as the "intelligent agents" which sense their environment and take some action to maximize their chance of success with respect to some pre-defined goal. Thus an intelligent agent driving the car would perceive the surroundings using computer vision techniques and then decide in which direction to move or when to apply brakes. Similarly, an intelligent agent conversing with humans would decode the human speech, parse it, extract semantics from it and then reply accordingly. This mating of AI with probability and decision theory along with the immense computational power available today has enabled AI to regain its popularity and it is now rightly considered as one of the most important fields in tech world.

III. AI in Our Daily Lives

a) AI in Social Media

There are tons of raw data available at social media platforms and AI is now used extensively to make sense of that data. Using Machine Learning and Data Science techniques and coupling them with AI, social media platforms are now improving the everyday user experience. Facebook uses AI technology to automatically tag the photos, filter news feeds and figure out trending topics. LinkedIn acquired Bright in 2014 – an AI and Machine Learning based start-up – to offer better job-candidate matches for both potential employers and job-seekers. It uses Machine Learning algorithms to do this prediction taking into account the past hiring trends, job location, work experience etc. Similarly Pinterest has recently acquired Kosei, a data software company which specialized in personalized recommendation modeling. The motive is that using such technology would help them in recommending products based on content pinned on network. These are but few of the many examples prevalent in our social media networks and many of us are not even aware that many times it's AI which is choosing our next best friend or our next favorite product.

b) Search Engines

Many search engines have started to incorporate AI in their search algorithms to refine and improve the search results. Google is obviously leading the trend here. Google bought the British AI startup DeepMind in 2014 at the whopping $400 million to kick-start its AI ventures. Since then it has attracted a number of leading AI researchers from both academia and industry who are doing cutting-edge research in various AI domains. Google has recently incorporated RankBrain [13] in its overall search algorithm Hummingbird. RankBrain is an AI based system which helps the main algorithm in processing search results. Just like famous PageRank [14] – a ranking algorithm for ranking search results, RankBrain also helps the main algorithm in processing of results and refining the user search queries. Google uses over 200 "signals" to rank the search results and PageRank is the most important signal but RankBrain is also not far behind

12

and is currently the third most important signal. Other search engines have also started to use AI and Machine Learning to provide more targeted and more refined search results to their users.

c) Strategy Games

Defeating humans in computer games was perhaps the first biggest achievement of the AI. Even though strategy-based computer games have been quite popular for some time now, it is really in the last decade that AI has got nearly invincible. For turn-based games like Chess and Checkers, and other sports games like Soccer, Baseball and Tennis, if AI could be allowed to play to its full potential, there is literally no chance for humans to win. So much of the effort these days goes into making the gameplay more realistic and letting the humans win eventually if they play good enough. In fact this is one of the biggest difference between the research community and industry in AI games arena – while researchers tend to make AI better and stronger so that it can't be beaten even by the best of human players, industry tends to focus on real-time and real-world experiences and tunes AI in a way which can keep the human opponent engaged and not make the AI seem invincible. Early games made the mistake of making AI too strong, one example of which was *Empire Earth* which had wonderful gameplay but the AI at its strongest was impossible to beat, even by the best of players since it could collect resources at an alarming rate and build forces in no time. More famous games like *Age of Empires* and *Command and Conquer* more or less got the AI part right and players with enough experience were able to outmatch the AI. Recently, as computers have become much more powerful and nearly every PC comes with GPU(either in the form of integrated graphics or external graphics card), strategy games now tend to focus more on graphics part of the game to make the objects and animations look real. Moreover, with the current processing capabilities at hand, it is harder not to let AI grow stronger than before. But the game makers have to strike a fine balance where AI is not so weak that it can be outmatched pretty easily but also not so strong that no one can beat it. Besides, AI now tends to be unpredictable so that a cunning and perceptive human can't detect the patterns on which it operates and make strategy accordingly. So every time you play against AI, even under same scenario, AI tends to mix things up just like humans do. AI role in games is certainly one of the most prevalent one in our lives, especially in the lives of teenagers and perhaps the source of the fear for some AI skeptics who believe that one day AI can easily beat humans in the real world just like it can easily beat humans now in virtual world.

d) Self-Driving Cars

One of the most promising application of AI in near future would be the self-driving or autonomous cars. Using advanced machine learning algorithms, these cars would be able to navigate through highly crowded and busy roads and could run on many different kinds of terrains. There is already a huge progress made in this area with some big names like Google, Tesla and Uber investing big-time in self-driving cars. Some cars already have self-driving features in them in which a human driver can turn the auto-drive mode on but they can be overridden by human intervention, pretty much like cruise control. An actual self-driving car would not require any human intervention and it would navigate using its sensors and radars. As per Google [15], a self-driving car would be continuously answering these four questions:

1. *Where am I?* - The car would use map and sensor information to determine where it is at any given moment.
2. *What's around me?* - The car would detect objects around it using sensors and cameras and classify them according to shape and size. Google's self-driving cars can detect objects from as far as two football fields.
3. *What will happen next?* - The software installed in car then decides what is going to happen next? Which object will move, which will remain static?
4. *What to do?* - Finally the car decides what to do next. Do brakes need to be applied or not? Is it safe to accelerate the car?

Google is the leading researcher in this field with its cars having gone the testing of over 1.5 million miles. It previously used customized models of Lexus but since 2014 it has been testing on its own specially built prototype. Recently Uber has launched its self-driving fleet in Pittsburgh which for the time-being will also be monitored by human drivers. Similarly, Tesla has been providing auto-pilot feature in some of its cars for some time now. Some states like California, Nevada and Texas have already passed legislation regarding self-driving cars while others are contemplating doing so. It is everybody's guess when self-driving cars would take over the human-driven ones but with the recent progress, that day doesn't look too far away.

e) AI in Military

Considering that much of initial funding in AI came from DARPA, it is no surprise that AI is pretty heavily used in military and warfare these days. Unmanned Aerial Vehicles (commonly known as drones) and Unmanned Ground Vehicles (UGV) have been in use of military for over a decade now. Famous among them are the Gladiator Tactical UGV (used by US Marine Corps), ViPer (used by Israeli Forces), Sarge and The Warrior (Unmanned Tanks) and The Talon (used for bomb-disposal). Similarly drones have been used by US for bombing militant hideouts in Afghanistan and Yemen. More recently, domestic law enforcement agencies have also started using AI robots for bomb-

13

disposal missions. In fact, very recently, the killing of Dallas shooting suspect through robot is widely believed to be first such incident [16] where a bomb-disposal robot was itself armed with remote-controlled bomb and detonated when it went near the suspect. Use of AI in military is a grey area and the use of drones and other unmanned ground vehicles have been large criticized by human-rights organizations.

f) Speech Recognition & Personal Digital Assistants

Ever since HAL-9000 made its debut in Stanley Kubrick's famous *2001: A Space Odyssey*, people only perceive of AI as a talking machine; such was the cultural impact of that movie. But ironically this is perhaps the trickiest and the least robust of all AI applications. The notion of having an AI talking with you and doing your several tasks is no more a dream now but many challenges still loom. Most significant among them is correctly recognizing the human speech and almost equally challenging is the task of making sense out of this speech. The conversational bots have become quite common but many of them are text-based and domain specific. On smartphone side, Google Now, Siri and Cortana are pretty state of the art and can do almost everything short of having a full-fledged conversation with you. They can make calls, send emails, tell weather, recite important news and many more. Much of this has only been possible recently due to advancements in processors and memory. Using deep learning, extremely sophisticated speech models can be built and custom tuned to the voice of smartphone user. Such dialog systems are extremely useful for people who are less tech savvy as they can just order their phone to do things for them rather than navigating the phone for the desired functionality. As more and more powerful models are being built, use of such systems is becoming more ubiquitous. Focus now is on developing such systems for local languages so that people unfamiliar with English can also have a taste of it.

g) Recommender System

Recommender Systems are now used by all digital marketing vendors and even blogs and social websites. The idea behind them is to observe the patterns of the certain user and then make recommendations to the user based on past behavior. For instance, you shop on Amazon and it will give you a list of recommended items. Similarly, you watch some TV shows on Netflix or Hulu and they can make recommendations to you based on your interests. This trend of targeting users individually and making recommendations to them based on their behavior is a huge plus for marketing people. And behind all of this is sophisticated AI primarily based on unsupervised machine learning algorithms which mine for patterns in the user behavior and draw conclusions accordingly. Websites also have now started to post more directed ads to users based on their browsing behavior. Some believe it to be an invasion of privacy, but this is the price we have to pay for living in the digital world – nothing is a secret anymore and businesses tend to exploit it to their advantage.

h) Robotics

No discussion of AI can be complete without mentioning robots – the physical manifestation of AI. While most other AI products can operate on simple general-purpose computers, robots require special hardware and a wide array of sensors to operate in a seamless manner. For most part of AI history, robots were of little practical use and much of the work was done by hobbyists and AI enthusiasts. In last decade or so, research community picked up the Robotics fever and started organizing competitions and contests like RoboCup [17] which pitted robots against each other in various contests and the winner was awarded sizeable reward. Owing to the recent advancements in processing and memory capabilities, and the availability of very high precision sensors, robots are now used in industrial applications as well. They are used to perform high-precision jobs like welding and riveting, used for material handling and assembling the products, used in ultra-high precision surgeries and also used in potentially dangerous situations like toxic-waste cleaning and bomb disposal. Japan is the leader in designing and making highly advanced humanoid robots, most famous of which is ASIMO (Advanced Step in Innovative Mobility) [18]. It can walk and run on smooth as well as uneven or slippery surfaces, climb stairs and pick and drop objects. It can also recognize human commands and human faces and can avoid obstacles. Similarly NAO [19] is another famous robot which can act as "true family companion" for families. But undoubtedly, the world leader of robots is NASA's Curiosity Rover [20] which has been exploring Mars since 2012 and has sent some amazing pictures of Martian terrain back to Earth. Its primary mission is to determine the Mars habitability and search for any potential life-forms like microbes. The way robotics industry is progressing, it won't be longer than two or three decades when robots would become ubiquitous in every household for doing simple everyday chores like washing and cleaning.

IV. Concerns Regarding AI

While AI has garnered considerable support over the last half-century or so, the recent advances have made some people afraid of its potential strengths and misusage. These concerns can be broadly categorized into 3 main areas:

a) Controversy regarding "Rise of Machines"

This notion that one day AI is going to take over and make humans their slaves or even worse, make humans extinct, is not new. Some of this fear is fueled

by sci-fi novels and movies like Terminator but lately some of the big names of modern science have also expressed-concerns over unbridled progress of AI. Most notable among them are Stephen Hawking, Bill Gates and Elon Musk. The controversy regarding Artificial Intelligence or "super-intelligence" has been fueled by Oxford University philosopher Nick Bostrom [21] in his articles and books where he presents several hypothetical scenarios in which AI takes over humans. Stephen Hawking also pitched in with his two cents, hypothesizing that one day AI would become so powerful that it can create a better replica of itself and it would set a chain of better AIs which would eventually no longer need humans and would be dangerous for humans. Similar concerns have been shown by Bill Gates and Elon Musk [22]. But in view of most of the AI community, these concerns are far-fetched and perhaps too distant. AI technology would probably take hundreds of years more to reach at any discernable dangerous level for humans. According to one AI researcher, worrying about AI taking over the world is analogous to worrying about over-population on Mars.

b) Ethical & Moral Concerns

The more serious, and legitimate, concerns are raised from moral and ethical points of view. Is it right to give military AI robots the power to kill the enemies or decide their fate in some other way? Is it right for AI recommender systems to display ads to people based on its perception of people? Is it right to make AI moral agents – i.e. give them the power to decide what is right or wrong in a given scenario? All of these questions deliberate on the fact that how much power can be entrusted to AI! Moreover, most modern AI systems are constantly evolving and learning from their interactions. There is an inherent danger that they would learn to mirror the human values and those values would be biased, based on the type of people AI would interact with. One such example we have already seen in the form of Microsoft Twitter Bot, which learnt profanity and racism pretty quickly due to its interaction with people who were deliberately trying to misguide the bot and were eventually successful [23]. Thus there would always be possibility of AI becoming a representative of ideas and values of small group of people rather than human population as a whole. There is also a question mark on the power of people behind programming AI and their ability to program AI in a negative way can be disastrous. There is also concerns about hacking of AI related products as recently demonstrated in DEF CON 24 [24]. The hackers were able to take control of an autonomous car and were able to accelerate it and apply brakes. Since AI is pretty heavily dependent on sensors, it is also a potential area of concern as some hackers were able to blind the Tesla auto-pilot car and make it collide with objects. All these concerns are pretty legitimate and there is a need to make AI more

secure and avoid its transgression in realms of moral and ethical decisions. Professing this point of view, many among AI community have proposed a ban on use of AI in military endeavors.

c) Financial & Social Concerns

There are also financial concerns at stake with AI permeating more and more areas of human life. Self-driving cars, when common, would pretty much make the human drivers obsolete as many of traffic accidents are due to human error. Also it would be more financially feasible for companies to have self-driving trucks rather than a man behind the wheel. Similarly AI when advanced enough to have a natural conversation with humans, would replace the service center receptionists, especially those behind the phones. In the industry too, AI robots are becoming advanced enough to replace human jobs of packaging and assembling but that would take some time. Amazon recently held a competition for the fastest robotic assembler which could categorize and arrange objects correctly and quickly. But many AI researchers have brushed-off these concerns as typical conservative response to anything new. They cite the example of ATM machines which are pretty common these days but were faced with a lot of criticism when launched as they were supposed to make bank staff members out of job. But introduction of ATMs opened newer avenues of interest for the human staff. In a similar way, while AI could replace some of the human jobs, it can open many more opportunities for the humans.

V. CONCLUSION

There is no doubt that AI has become a major part of our life now, and for better or worse, it is bound to remain an integral part in future. It is already playing an important role in several domains like personal digital assistants, recommender systems, autonomous cars, social media and many more. In coming decades, AI is likely to grow even more and become even stronger. This fact has made some people wary of its success and they are suggesting to put a lid on its progress to keep it under control. Most of their fears are unfounded but some are legitimate and needs to be addressed.

REFERENCES RÉFÉRENCES REFERENCIAS

1. Turing, A.M. (1950). Computing machinery and intelligence. Mind, 59, 433-460.
2. Russell, Stuart J.; Norvig, Peter (2003), Artificial Intelligence: A Modern Approach (2nd ed.), Upper Saddle River, New Jersey: Prentice Hall, ISBN 0-13-790395-2.
3. Mc Corduck, Pamela (2004), Machines Who Think (2nd ed.), Natick, MA: A. K. Peters, Ltd., ISBN 1-56881-205-1, OCLC 52197627.

4. Simon, H. A. (1965), The Shape of Automation for Men and Management, New York: Harper & Row.

5. Natural Language Input for a Computer Problem Solving System, Bobrow's PhD thesis-http://dspace. mit.edu/handle/1721.1/5922.

6. Crevier, Daniel (1993), AI: The Tumultuous Search for Artificial Intelligence, New York, NY: BasicBooks, ISBN 0-465-02997-3.

7. John McCarthy. "LISP prehistory - Summer 1956 through Summer 1958".

8. Arthur E. Bryson, "A gradient method for optimizing multi-stage allocation processes", April, 1961. In Proceedings of the Harvard Univ. Symposium on digital computers and their applications.

9. http://www.ibm.com/ibm/history/ibm100/us/en/icons /deepblue/.

10. Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., Fong, P., Gale, J., Halpenny, M., Hoffmann, G., Lau, K., Oakley, C., Palatucci, M., Pratt, V., Stang, P., Strohband, S., Dupont, C., Jendrossek, L.- E., Koelen, C., Markey, C., Rummel, C., van Niekerk, J., Jensen, E., Alessandrini, P., Bradski, G., Davies, B., Ettinger, S., Kaehler, A., Nefian, A., & Mahoney, P. (2006). Stanley: The robot that won the DARPA Grand Challenge. Journal of Field Robotics, 23(9), 661–692.

11. http://www.ibm.com/watson/

12. "Research Blog: AlphaGo: Mastering the ancient game of Go with Machine Learning". *Google Research Blog*. 27 January 2016.

13. Clark, Jack. "Google Turning Its Lucrative Web Search Over to AI Machines". *Bloomberg Business*. Bloomberg. Retrieved 28 October2015.

14. Brin, S., Page, L. (1998). "The anatomy of a large-scale hypertextual Web search engine" (PDF). Computer Networks and ISDN Systems. 30: 107–117.

15. https://www.google.com/selfdrivingcar/how/

16. https://www.theguardian.com/technology/2016/jul/0 8/police-bomb-robot-explosive-killed-suspect-dallas

17. Kitano, H., Asada, M., Osawa, E., Noda, I., Kuniyoshi, Y., Matsubara, H., "RoboCup: The Robot World Cup Initiative", Proc. of the First International Conference on Autonomous Agent (Agent-97), 1997.

18. http://asimo.honda.com/

19. "Unveiling of NAO Evolution: a stronger robot and a more comprehensive operating system". Aldebaran Robotics. 2014.

20. https://www.nasa.gov/mission_pages/msl/index.html

21. Bostrom, Nick (March 2002). "Existential Risks". *Journal of Evolution and Technology*.

22. "The Future of Life Institute Open Letter". The Future of Life Institute. Retrieved 4 March 2015.

23. *Wakefield, Jane. "Microsoft chatbot is taught to swear on Twitter". BBC News.*

24. https://www.defcon.org/html/defcon-24/dc-24-index.html.