

John Zhang *, Ming Fan, Bin Gu, Vijay Mookerjee, Bin Zhang, J. Leon Zhao

Business Values/Implications of AI and Machine Learning

<https://doi.org/10.2478/dim-2018-0016>
received July 27, 2018; accepted September 30, 2018.

Keywords: artificial intelligence, machine learning, business value, deep learning, self-learning, information system, decision-making

1 Introduction

Artificial intelligence (AI) and machine learning are the major driving forces of the Fourth Industrial Revolution. These technologies bring new methods and solutions to many disciplines under the umbrella of Information Systems (IS). With the help of AI and machine learning, the IS discipline is able to study more emerging phenomena in information technologies and process big data. At the 12th China Summer Workshop on Information Management (CSWIM 2018) in Qingdao, China, the meeting's chairs assembled a panel of leading IS scholars who work intensively in AI and machine learning to discuss the theme "Business Values/Implications of AI and Machine Learning." The panel consisted of Ming Fan, Bin Gu, Vijay Mookerjee, Bin Zhang, and Leon Zhao. John Zhang was the panel coordinator, who asked the panelists to discuss about the following topics: 1. Business use cases of AI or deep learning; 2. Potential interesting phenomenon, research opportunities, decision problems, and challenges in the era of new technology, big data, and AI; 3. Implications of AI and machine learning for the IS

field. This paper is a narrative of the panelists' speeches and the ensuing discussion.

2 Ming Fan: "AI is pushing the boundary of what computers can do"

The panel today will discuss Artificial Intelligence (AI). It is a continuation of the conversation regarding the study of AI in the IS field. There are a lot of things going on in the IS field. AI will be one of the most exciting parts of this field in the future. We are going to talk about what the business cases are and the possible implications for research in the IS field. I have some exciting business examples of AI and deep learning to share. Last week, I happened to see a video about Google's conversation agent. Google's CEO, Sundar Pichai, demonstrated an application in which the AI agent made a reservation at a restaurant and started a very smooth conversation with a human. It understood human language and the complex situation in which it found itself. On the other side of the conversation, the person from the restaurant could not detect that he was talking to an AI machine. That is a very impressive application.

When I view AI's applications from a researcher's point of view, I wonder what they mean for us. I have done some research in this area. Let us consider how the use of computers impacts the labor market. According to economic theory, the use of computers could replace routine human tasks. For example, computers are really good at storing data, inputting data, and other repetitive work. This is routine work. But it is difficult to replace humans for non-routine jobs. If you look at what AI is doing now, you will find that AI is pushing the boundary of what jobs computers can do. AI, through deep learning, is starting to replace humans for some non-routine tasks. However, I have some concerns.

I think AI is very promising, and it is developing so fast. I think about the unlimited power of computers and the capabilities of computers in terms of learning. This reminds me of another theory, or law, that we talked

*Corresponding author: **John Zhang:** Department of Information System, W.P. Carey School of Business, Arizona State University, Tempe, AZ 85287, USA, E-mail: Zhongju.Zhang@asu.edu

Ming Fan: University of Washington, Foster School of Business, Seattle, WA 98195, USA

Bin Gu: Department of Information System, W.P. Carey School of Business, Arizona State University, Tempe, AZ 85287, USA

Vijay Mookerjee: Naveen Jindal School of Management, University of Texas at Dallas, Richardson, TX 75080, USA

Bin Zhang: Eller College of Management, University of Arizona, Tucson, AZ 85721, USA

J. Leon Zhao: Department of Information System, College of Business, City University of Hong Kong, Hong Kong, China

about earlier in the IS area: Moore's Law. I heard about this theory last year and I was really excited. There is also a theorem called the universal approximation theorem. Basically, the idea is proven in theoretical computer science. It states that a feedforward network with a single hidden layer containing a finite number of neurons can approximate continuous functions on compact subsets of \mathbb{R}^n . The power of a computer with enough data and the ability to learn is unlimited.

Let us return to the concerns I mentioned earlier. When AI agents start to replace certain human interactions, such as booking tickets, making restaurant reservations, etc., the question I have is, with the development of AI, what is left for us human beings? Are we becoming less and less engaged? I have also read some other books. It is well established that social interactions are extremely important in human brain development and human evolution. But AI agents make us less engaged and less interactive. So, my concern is that it may be bad for human brain development and evolution.

Another issue I have, as I mentioned earlier, is the unlimited power of learning. With lots of data, AI can basically learn how to make more accurate predictions. The impression I have is that there are a lot of processes and much decision-making in the real world that are imperfect. For instance, we have a lot of biases in law enforcement and college admission. If we use AI technologies in all these areas, any biases would be amplified and reinforced by the AI algorithms. That is another issue I think about.

3 Bin Gu: “Characteristics of AI: Self-learning, non-linear development, and new biases”

Thanks, Ming. The human race is facing an exciting, albeit somewhat uncertain, future in the presence of AI. When John first asked me to serve on this panel, I told him that I have not done much research on the technical side of AI. He suggested that I discuss the big issues concerning AI, especially those with economic and societal impacts. That is what I would like to address today.

When we review the history of IS research, we have a substantial amount of research examining the role of IT and its impact on individuals, organizations, and societies. For example, Lorin Hitt (Wu, Hitt, & Lou, 2017; Wu, Jin, & Hitt, 2014) and Erik Brynjolfsson (Brynjolfsson, Rock, & Syverson, 2017) have undertaken a series of studies on how IT investments affect firm performance and organizational

structures. As IT has evolved from mainframes to PCs to big data, IS research has naturally evolved to study the impact of these new generations of IT. Now we are talking about AI. Is AI just another IT application or an advanced version of previous IT initiatives? What is unique about AI? As an author and researcher, I really want to know what is really new about AI. After many conversations, I have come to the conclusion that, compared with the previous generations of IT, AI has three unique characteristics.

First of all, AI is self-learning. One thing I believe very different from the previous generations of IT is that AI can learn from itself. Through self-learning, it improves continuously. This makes AI quite different from the standard economic approach of viewing IT as a static input. This is something unique in my opinion and may deserve our consideration from a research perspective.

Second, when something can learn by itself, it leads to non-linearity, which means the more you learn, the faster you learn. Everyone here has heard of the term singularity. It is the point in time when AI has learnt so much that it starts to surpass the human race, which means that it no longer needs to learn anything from humans. I am not predicting that this singularity is coming. But, I do think that non-linearity could be something unique that is worth our attention.

The third characteristic is related to how AI affects human decision-making. It is an interesting question that has intrigued researchers for decades. In the 1950s, Brunswik developed a decision-science model called the Lens model. A series of experiments based on the model showed that human beings are not Bayesian learners and that computers can outperform human beings by learning human decision variables while eliminating decision biases. Researchers have carried out many experiments to confirm this. These studies were conducted long before we had super computers and fast machine-learning tools. Therefore, I think the final interesting question concerns if, when, and how well, machines can learn human decision variables while eliminating the critical mistakes that humans would make. Although AI can remove some decision biases, it does not necessarily mean that AI can remove all biases. Indeed, as we have seen in many cases, AI may pick up race or gender as quality signals and thus amplify race or gender biases. More broadly, there is a need to teach AI ethics.

In summary, my points are that AI has three unique characteristics that distinguish it from previous generations of IT: AI can learn by itself; AI is non-linear in its development and growth; and AI may reduce decision biases while amplifying other forms of biases.

4 Vijay Mookerjee: “The development and implications of AI”

We have not talked much yet about how AI came to be the way it is today, which could be helpful for prospective research. To the best of my knowledge, people have been talking about AI since the 1950s. Let us consider Alan Turing, who is considered the father of computing, and John McCarthy at MIT and Marvin Minsky, who developed complicated thinking machines. When the three of them got together and started talking about AI, Alan Turing proposed what was called the “imitation game;” there was also a movie about this. The imitation game essentially entailed a human dealing with an agent that he/she could not see. The human talked to the agent and then had to decide whether it was a human or not. If a computer could fool a human into thinking that he/she was talking to another human, it won the “imitation game.” At that time, McCarthy predicted that computers would be able to win the imitation game in 2000.

Today, we have mostly considered AI in specific contexts, such as booking a hotel or a restaurant. But Turing considered AI in a more general sense – in any situation. Can an agent fool a human? If so, is it “human?” He realized that computers had not yet reached this point. In the 1960s and 1970s, what was really fascinating was game playing, such as checkers, chess, and similar games. Very soon, it was realized that using general purpose methods to solve complex work would not suffice. It entailed the paradox of computers needing lots of domain-specific knowledge. That is why we developed expert systems.

You may have heard of some of the most famous expert systems, e.g. Mycin and Dendral, developed at Stanford, which have been used to diagnose bacterial infections. If you examine the first phase of AI in the 1980s, you may feel a little disillusioned for the AI researchers. Remember, at that time, AI was not a hot topic yet. In retrospect, what had happened was that the scope of AI had begun to be narrowed. A great deal of conceptual knowledge was needed to solve AI problems at that time. However, things became very different in 1980s. We have just mentioned Moore’s Law, which is still valid today. Thus, computers are becoming twice as powerful every 18 months. The benefit of this for AI is that fewer computers are needed to facilitate distributed computing for machine learning.

We have also talked today about big data. The presence of lots of data has enabled the start of machine learning and this has become a very important part of AI.

Significantly, large corporations have now started to invest in AI, such as Google and IBM. It seems that many of their products have been successful. Recently, Google’s program beat the world chess champion. Following this, Google released another version in 2017 by making the program learn from playing itself. After playing against itself 5 million times, it created a new version called *AlphaGo Zero*. Subsequently, *AlphaGo Zero* beat AlphaGo by 100 games to 0. The progress is obvious. In 2015, the first program beat the world champion, and, two years later, a program was produced through enforced learning that beat the original version by 100 games to 0. This shows how fast the technology can advance. Therefore, I think that we should leave behind the disillusionment mentioned previously and consider instead the promising and optimistic situation that has arisen in the last 20–25 years.

It is quite clear that anything humans can do is based on processing and it appears that, eventually, computers will be able to not only equal us but also surpass us. There are many examples of this. Computers are extremely good at vision and speech recognition. You can query your bank account using your voice and the computer can recognize who you are and authorize you. Speech recognition is important. As I have already mentioned, in game playing, computers have beaten humans. Diagnosis is another interesting example that has recently used deep learning. Computers examined about 700,000 records of people who had been to the doctors and proved to be much better at predicting diseases than the doctors. The problem is that, although the program outperformed doctors, doctors were not happy about adopting it. Because they did not know how it worked, they were unwilling to adopt it. But in areas such as decision-making, examining medical records and recognizing fast-moving objects rapid progress is being made. For example, in China, video cameras can take a picture of a passing car and get a very clear picture of the driver. A policeman would not be able to see this because it happens so fast, but the camera can. These are the areas in which we have to admit that, eventually, computers will outperform us.

From the business perspective, I was thinking about the obvious issue of job replacement. A good analogy for this is outsourcing. Manufacturing outsourcing began in the US in the 1980s. Subsequently, in the 1990s, service outsourcing began. Initial projections were for job losses of approximately 50%. However, what we have seen is about 80% of the jobs being retained. This has happened because of global tensions and humans’ resistance. In essence, when humans encounter a technology that threatens them, they will find ways of not adopting it;

they will find ways to block it. Today, the projection is that approximately 40% of the knowledge and jobs will be replaced by AI in the next 15 or 20 years. Personally, I do not think this will happen so dramatically. It will not happen as quickly as predicted if humans can be more effective than the computing replacement.

There is one last thing that I would like to touch upon briefly. If you think fundamentally, our brains process using memory. Any task that can currently be so reduced, the computers will be able to perform. My belief is that all intelligence is fundamentally not able to be replicated by machines because intelligence is conscious. We do not need memory to be conscious. Consciousness is a memory-free ability. Machines can never have consciousness. They are able to automate certain things but human intelligence, because of consciousness, will always be a few steps ahead. What might happen is that, whether you view it positively or negatively, we might get a lot more free time. I am an optimist. I feel that AI systems will lack certain human qualities.

However, this means that we need to find ways of spending our extra free time. For example, advertising and entertainment would become very important. Because we will be sitting in front of various entertainment devices to spend our extra free time, I think there will be a big boost for the advertising industry. I think that this is one of the areas that IS people could look at.

Another area I believe we should concentrate on is coordination problems. AI is no longer a singularity. It has spread across IT systems. This obviously also involves various parts of the Internet starting to talk to each other. The important questions will concern how best to coordinate this and how much information to exchange. For those who remember the book from the 1970s called the *Economic Theory of Teams* (Marschak & Radner, 1972), I think team theory will again be important. The book was way ahead its time. But there was no real application at the time this book was written. But today, I think there is a lot of potential for these ideas. So, I urge you to think about the problems that might not be business problems, but very interesting. Think about what is all around us in the ether – information.

Consider also the drones that can now sense each other while they perform their back flips or complete their circuits. Think then about a more complex situation like surveillance. The drones can collect data at different points and then send the information to each other or may be to a hub. These sensory data will reveal anything wrong. So, cooperating AI, or distributed AI, will definitely become important in big teams.

Another thing that I believe that will become very important is the extension of ID security through AI. Think about how you feel about AI. You may think AI is malicious. Let us say that you want to take a self-driving car to a 7-Eleven. If you input incorrect time data, you will be successful if you learn from your mistake. Similarly, ID security will be important because it might have more serious effects. For example, I want to watch the movie *Samba* because it is a highly rated one. But a friend tells me it is not a good movie, despite that the reviews says it is good. I would believe in him rather than the rating. I believe that AI systems have started to catch up on this. There may be a chance that AI is making probably malicious broadcasting in the network. Leon [Zhao] has been thinking about blockchain and perhaps blockchain might lead the ID security of distributed AI.

5 Bin Zhang: “How we can use AI for IS research and potential topics”

Thanks to the three previous extraordinary speakers for having addressed the impact of AI on IS at a macro level. I would like to address this topic at a relatively micro level. I want to talk about the algorithms that we could use for problems related to our research, potential topics we can study using AI, and the future of this domain.

Before I start, I would like to share with you a story. Back in mid-2000s, many computer scientists considered neural networks as algorithms that could achieve very good performance. But since this was a black box, it had significant computational cost problems. It was not considered as a solution to large-scale data, and thus, neural networks did not seem to have a future. At that time, algorithms such as Bayesian networks were considered as a preferred solution. But look at how popular and powerful deep learning is today. The point of my story is that we probably are not correct all the time, especially myself. So, what I say today is only a reference for you.

Currently, I think AI could be considered as one of the driving forces for the Fourth Industrial Revolution, through which we would like to connect engineering, biology, and computer science. Within AI, the method at the vanguard is deep learning. It will continuously be the leading algorithm development for the next 5 or 10 years.

Deep learning is basically a neural network with more than one hidden layer. Currently, there are two families of deep-learning methods frequently used by the IS discipline. The first is Recurrent Neural Networks (RNNs), which have applications generally in speech recognition

and content analysis. I think there is still a lot of space for using RNNs. Particularly, there is one type of RNN, long short-term memory method (LSTM), that we could be used to study many IS phenomena. Consider social media; we can consider social media as a very large data repository. In the earlier days, if you would like to know people's opinions or how they perceived a particular product, you had to conduct a survey and ask them. Nowadays, nearly everyone is using social media to record their life and thoughts. Therefore, what they think is recorded on the Internet. Seemingly, individuals are less sensitive about privacy. I have undertaken previous research on this topic and found empirical evidence. People tend to be more open and more willing to talk about their private preferences on social media than in the real world. Individuals tend to share their lives and opinions about products online, which provides us with many opportunities to analyze such content.

In general, content analysis in IS research is still primitive. Most of us are still doing sentiment analysis for a given post or paragraph to determine whether this user has a positive, negative, or neutral sentiment. We use the sentiment as a proxy to evaluate individuals' ideas or opinions. But this is relatively inaccurate. A positive sentiment in a post does not necessarily mean this person supports the agenda. What we should do is to use a tool like LSTM to analyze the meaning or extract the knowledge from the content. Then, we can determine individuals' opinions about a topic more accurately. We have a lot of content online right now. For example, in the online health community where doctors help patients or patients help each other, people never think about how good the quality of online medical knowledge is. Misleading information may have a negative impact on patients' health. We can use LSTM or, even better, bi-directional LSTM to extract the knowledge from the content. In some domains in particular, knowledge is presented in certain semantic formats. Let us take medical science as an example and assume that we have lexicon containing words in the medical science domain. We can define medical knowledge containing causality about treatment and disease. Assume that the first word in a sentence is found to come from the lexicon about medicine or is a method for treatment, the next word is "treats." Then, the third word is a disease from the lexicon. One can be fairly certain that this sentence is medical knowledge. We are able to use this method to evaluate the quality of the content online. If user-generated content contains a paragraph or many sentences with this format of medical knowledge, we can claim that the content is of high quality and very useful to patients.

Another useful deep-learning model for IS problems is Convolutional Neural Network (CNN). This model is frequently applied to image recognition and video mining; it can extract objects from images and videos. Image processing or computer vision is a very important component in robotics. If you want to develop a robot that can walk or you want to develop a self-driving car, one of the most significant jobs in robotics is video processing. A self-driving car mainly collects data and knows its environment through real-time video. A car needs to process and analyze real-time video, identify all the objects in the video, and then determine what to do. Assume that we use a deep-learning model to solve such a problem; our output could be turning left 30 degrees, turning right 90 degrees, decreasing speed by 20 miles per hour, or completely stopping. For a self-driving car, we need to solve the video-mining or image-recognition problems first. The rest of the problems, such as mechanical control, would then be easy. The reason why I bring up self-driving cars is to emphasize the importance and imminence of using unstructured data, including images and videos. There are a lot of unstructured data on social media. So far, we still intensively rely on structured data, such as numeric values, or fixed-format text, to study the problems we are interested in. We still have not leveraged unstructured text very much, not to mention audio, pictures, or videos. Recently, there was a study by researchers from Carnegie Mellon University on whether the quality of pictures on Airbnb could increase the property owners' revenue. In short, they found that, if the property owner provided a beautiful picture of the room, it could be rented at a higher price. The owner thus could make more revenue. There are so much unstructured data on social media, such as long texts, audio, pictures, and videos. If we were able to use CNN techniques and extract objects from the picture and possibly infer the context of the story, this could significantly complement the existing analysis.

Next, I would like to talk about some IS problems that can potentially be analyzed by deep-learning methods. Some of them are actually what I am working on. First, one phenomenon for which we can use deep learning to study is the social media problem, as I have mentioned earlier. That is the topic I am currently working on. On a video-sharing social media, for example YouTube, or Youku in China, there are many patients, especially chronic-disease patients, searching for treatments for their diseases, for example, insulin injection or physical therapy. Quite often, there are many videos about a disease and the quality can be heterogeneous. Such problems may be more severe in China. When searching videos about health

care or medical treatment, we see that many videos are actually advertisements. These videos contain low levels of medical knowledge and are not helpful to patients. In order to find videos that contain high-quality medical knowledge and are more helpful to patients, I analyzed both the videos and the metadata about the videos, fit them into deep-learning algorithms, and evaluated the quality of the knowledge contained in the video. I used natural language-processing algorithms, such as RNN, to process text data, and then inferred the quality of the knowledge. At the same time, I also used CNN to analyze the video to infer whether the video contained contextual information about the knowledge. If they matched, it was very likely that it contained high-quality knowledge.

The second problem we can study concerning AI is its value for enterprises. Many leading companies have invested a lot of money in AI. In 2017, Wall Street spent 1.5 billion dollars on its AI infrastructures, and 40% of the traders and financial analysts lost their jobs. But is this a good idea? Enterprises invested a lot of money in AI in the early days, aiming to achieve long-term cost savings compared to human labor costs. So far, there have been no studies investigating whether the investment in AI infrastructures was really worth the money. Sometimes, a technology also requires continuous upgrades and enterprises need to keep investing more money. This is another problem we can look at.

Another important problem concerning AI that we may study is data annotation. In order to make deep-learning algorithms work, ideally, we need millions or tens of millions of training samples to achieve good prediction performance. Understanding how to annotate data accurately is an important issue. If we have a small amount of data, we could hire domain experts or coders to annotate the data. But when we have a very large amount of training data, manual annotation is nearly impossible. Annotating data accurately and quickly for big data is both practical and necessary. I hope our discipline can soon provide a solution.

The last problem concerning AI that I suggest we study is the causality of the results generated by deep-learning methods. Since the deep-learning method is a black box, a research has to trust the results without any interpretation, otherwise one cannot accept the method. But in IS, identifying causality is very important. Although there are many open-minded editors in our discipline right now who appreciate research using deep-learning methods, we ultimately need to address this problem.

Finally, what does AI and machine learning mean for the IS field? The IS discipline is a very important component of the AI research community. Besides

designing AI algorithms, as performed by computer science, we also study the impact of such technology on society and enterprises. AI and machine learning provide a lot of opportunities for IS. They provide not only new methodologies for us but also phenomena for us to study.

6 Leon Zhao: “AI’s use of computational algorithms to replace human intelligence in decision-making”

Many people do not know that I am actually a “hidden” AI researcher. This is because I wrote my dissertation on expert database systems. My dissertation actually concerned the optimization of AI integration in databases. How can one materialize tables in the expert system? Why did I choose this topic? Part of the answer is that Silicon Valley posts hour-long videos with AI experts called *AI Show*. One time when I was driving and was trying to choose a dissertation topic, an expert on one show stated that, in the future, the issue would not be that we do not have enough data but that we would have too much data for people to assimilate. Therefore, AI would become dominant to analyze all the data to make intelligent decisions, which led me to the decision that I must do something related to AI. I continued along this path because, later, I found a job in a business school. I published four papers on data, algebraic algorithms, data engineering, and very large databases. With these four papers, I secured a tenured position in a small computer science department. But in business schools, these papers are not considered relevant. That is why I switched to workflows. So, I have become known as a workflow expert.

My definition of AI is quite simple. If you use computational algorithms to replace human intelligence in any context of decision-making, then that is AI. That is my definition. However, enough of the past; let us consider the present. I think the present situation is closely related to history. Vijay [Mookerjee] discussed the evolution of AI. When I was choosing my topic, I also considered voice recognition. In 1989 and 1990, I wrote several articles in this area. The best accuracy that voice recognition could achieve at that time was about 65%. But, 10 years later, everyone has started to use voice recognition. When you make a phone call, you are asked to say “one,” “two,” or three or “A,” “B,” or “C.” Once you have specialized the domain, the accuracy of AI is actually increased to almost 100 percent. If you just say simple

commands, how could that not be recognized? So, once you narrow down its application in special conditions, AI becomes very powerful. In fact, that is what has happened. Even AlphaGo is a very specific area. It uses deep learning and beats human experts. That is a lesson learned. If you are persistent, you will find ways to solve the problem; if you keep on digging, you will succeed. In the 1990s, there was a huge project in Texas to build a system comprising millions of rules. It turned out to be a big failure because of a lot of rules with numeric values. Those numeric values were based on empirical analysis. When conditions changed, those values became obsolete and so it was unable to operate. However, there are a lot of expert applications that have been successful. One of the famous examples is Campbell's Soup. When experts retire, they take with them the knowledge of how to make good soup. A successful expert system for soup making comprises about 50 rules.

Concerning the future, a lesson I have learnt myself is as follows: be persistent and you will always find ways to succeed. From my point of view, the most successful application in the future would be robotics. Robotics can be defined broadly; not just machines, but automatic arms and fetches. These are also robots. Nowadays, we have teaching robots that teach children English and we have agency staff and accounting robots.

Another area I have to mention is that I have just completed some blockchain research, which is obviously the frontier of IT. Not long ago, the President of China, President Xi, mentioned blockchain. Xi assigned blockchain the same rank as big data. After that speech, everyone started to talk about blockchain. Why is blockchain related to AI? I am sure that you have heard about smart contracts. Smart contracts are definitely closely related to AI. In fact, it can be called embedded AI – AI algorithms embedded in trusted data, which has a significant impact on business. In business, we make a lot of contracts. We have supply chains, supply-chain finance, and micropayments. In those environments, we can use embedded AI – the smart contract – to automatically execute business contracts. We can also use micropayments. Micropayments may not be efficient if everything is performed through humans because the processing cost is too high. Once you use embedded AI, or smart contracts, the transaction cost is very low, and so the size of the transaction can be small and you still make money. Imagine that, in the future, we will all be working on platforms: we will collaborate. Each of us will do something we are good at. Thousands of individuals will work together through e-commerce platforms, which replace some of the functions of big corporations. Smart

contracts would then be used to help people complete microeconomic transactions. My prediction is that, in the future, AI will somehow be embedded into transactional systems, changing the landscape for corporations and giving individuals more choices. Instead of going to work for companies, you could work for yourself with the help of platforms, such as open-ended workflow platforms. You could make contracts and deals using them and smart contracts will deliver your money; you will not need to work for corporations.

7 Questions and Answers

7.1 Q1 and Answer

Here, I have a question for Bin Gu. There are two views about AI. You ask a question and then the computer should learn. There is a paradigm referring to the main perspective that is called CASA, computer-as-social-actor. What do you think of the opposite view of AI, which is HMT, human-machine-trust? It is said that there might be a bias. We trust machines more than humans. That is called automation bias.

Bin Gu

If I understand it correctly, I think your question is about the opposite of injecting actions with machines because humans could have a bias toward trusting machine. I think it is true so far because humans do not think of machines as strategists. But I am sure this will change very quickly when humans realize that machines can be strategists. This is the simple answer to your question, and I think it is a really good question about what the difference is between machines and humans. Recently, my students and I started a project that wanted to investigate a very large company that had introduced AI into a warehouse. They developed a robotics tool for the warehouse but it turned out that the productivity had decreased. The question was simply why. We do not know the answer yet but, obviously, there are a lot of questions like this. How will machines work with humans and how will humans work with machines? What kind of trust do they have? In this case, one simple answer we have seen so far is that humans love to trust machines. As a result, they expend more effort trying to stop machines working. I think there are lots of interesting questions that can be studied in this interface.

7.2 Q2 and Answer

I have a question for Bin Zhang. You talked about high-quality knowledge. How can it be deciphered? I was just wondering if it could be the nature of the sentiment determined by these terms. For example, some people like applications while others like texts. What is the intention of a text? Is it aggressive, for example? Or is it not so aggressive? Can you give me a clear definition of sentiment?

Bin Zhang

If I understand it correctly, you are asking about whether we are able to set up a numeric value representing the degree of sentiment, not necessarily a binary value, is that right?

Q2 continued

No, a more detailed sentiment.

Bin Zhang

I see, so you are talking about more detailed labels about sentiment, instead of a simple binary classification, such as positive or negative. Sentiment can include sad, happy, angry, or arrogant. That is really a good question and a possible future direction for sentiment analysis. We could address this problem through higher granularity. Indeed, currently, our major results are still binary or triple, mostly positive and negative, sometimes neutral. But we do not have more detailed classification about sentiment, such as happy, sad, angry, arrogant, and so on. Such problems also relate to the annotation problem I mentioned earlier. The challenge is mainly in annotation. It is very hard to label data representing sentiment at such a level of granularity. It requires much higher requirements for coding with carefully labeled data. We can still solve such problems though. If we were able to find enough training data about different sentiments, accurately label it, and fit it to proper deep-learning algorithms, we would be able to classify data to multinomial sentiments. What you ask about could actually be a potential direction for sentiment analysis, which I have not worked on yet. The audiences sitting down there can consider it as a potential research topic.

7.3 Q3 and Answer

Since the topic concerns business value, I would like to hear about your experience or predictions about what you think about the business value of this domain.

Vijay Mookerjee

At this point, I do not think it is complete. Google, Facebook, and IBM have spent millions of dollars on it. If you look at it carefully, they have spent the same amount of money on more traditional investments over the last 15 years. For example, I think that, in the future, Watson will have a huge impact. It is Watson that essentially understands natural language and is able to go to many data sources to find the answers. Basically, Watson includes 95% of the answers in Wikipedia. That is what it does; it connects to Wikipedia. But, we can imagine that it could provide more detail than knowledge, i.e. intelligence and wisdom. Once that step is happening, it can begin to earn itself a living. But at this point, this is certainly futuristic.

Leon Zhao

JD.com in China just launched five delivery robots. They are now, in an experimental stage, running in some cities in China. They are making thousands of deliveries, replacing delivery men. Apparently, it is predicted that this will save a lot of money. In addition, they have obtained a license to deliver using unmanned airplanes. In fact, they have license to produce unmanned airplanes that carry 30 tons of goods. So, I think robotics equipped with AI will dramatically change the landscape of delivery employment.

Ming Fan

I think that is a very good example. Wall Street posts a series of videos called *Make Money*. The presenter is quite crazy but I like him. He talks about lots of technologies and their business value. According to him, there are two true AI companies in the US today. One of them is the Netflix. Netflix basically uses AI to learn the customer behaviors and watching behaviors to create content. He thinks what makes it a true AI company is the feedback loop and the virtuous loop. They get more data and create

better content for their customers. Customers pay very low prices. Because of the low-price model, it creates content that many customers watch, so they get more data, accruing billions of dollars every year. This could be the business value.

Bin Gu

I think that, looking at the history of all emerging technologies, some try to mimic humans while others do not. I think what we have seen in the AlphaZero case, as well as in many other cases, is that machines can do better: machines can replace humans. The AlphaZero case showed us that over hundreds of years, as human masters, what we had was heuristics. But, it was not the best solution. What my colleague at Austin researched a few years ago concerned traffic. Right now, we have many traffic lights on the road. In the future, there will be no need. They demonstrated that efficiency in the city would be improved by over 50% if we got rid of traffic lights, and all self-driving cars can rely on themselves. So, I think that is future.

Leon Zhao

This is a joke, so do not take it too seriously. Regarding the negative aspects of AI, it will reduce the importance of behavioral research in corporations. Why? Because humans will disappear. There will be robots running around the warehouses and delivering things on the road. JD's delivery robot is very small. It will not create traffic issues. Even if it bumps into a human, it will not kill the person. So that is the negative aspect of AI. It reduces the importance of behavioral research.

Vijay Mookerjee

Amazon used to have complex restocking systems. Humans were unable to recognize where to put things. So, they used heuristics to put things near to similar things. The whole thing was not being done by robots. Also, the room for a robot is very simple: find the shelves and put things in them.

References

- Brynjolfsson, E., Rock, D., & Syverson, C. (2017). Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics. *NBER Working Paper No. 24001*. Retrieved from <https://www.nber.org/papers/w24001>
- Marschak, J., & Radner, R. (1972). *Economic Theory of Teams*. New Haven: Yale University Press.
- Wu, L., Hitt, L. M., & Lou, B. (2017). Data Analytics Skills, Innovation and Firm Productivity. *The Wharton School Research Paper No. 86*. Retrieved from <https://ssrn.com/abstract=2744789>
- Wu, L., Jin, F., & Hitt, L. M. (2014). Are All Spillovers Created Equal? A Network Perspective on IT Labor Movements. *Management Science*, 64(7): 3168–3186.