

Report from IJCAI 2019, Top AI Conference for 50 years

Which Crisis is Coming First – of AI or of World Economy?

Matjaž Gams

Jožef Stefan Institute, Jamova 39, 1000 Ljubljana, Slovenia

E-mail: matjaz.gams@ijs.si

Editorial

1 Introduction

This year IJCAI [1] - "International Joint Conference on Artificial Intelligence" - celebrates half a century of continuous conferencing as the world AI's most important global event. Over the last ten years, the number of submissions has grown steadily, by more than 30% in the last two years alone, and for 2019 it approached 5000 with 2700 committee members (Figure 1). The acceptance rate this year was less than 18%, one of the lowest ever, leading to 850 papers presented at the main conference, accompanied by three days of workshops. The papers from China (327, 38%) toppled the papers from EU (152) and USA (169) combined, while other countries followed with Australia (37), India (20), Japan (18) and areas like Eastern Europe or Africa far behind. The shift in recent years has been enormous, not only in terms of the number of papers, but also in terms of the number of new applications and the overall focus of countries and resources involved. The conference costs for one speaker, for example, amounted to several thousand euros. In terms of rewards, lifetime achievements and invited lectures, USA and EU still dominated due to inertia since the number of senior AI researches in Asia has only recently started to increase.

Besides Program Committee Chair Sarit Kraus, there were Tutorials Chairs, Workshop Chairs, Demo Chairs, Doctoral Consortium Chair, Robot Exhibition Chairs, Video Competition Chairs, Survey Chairs, and Chairs for Sister Conference Best Papers, Journal Track, Special

Track on Understanding Intelligence and Human-level AI in the New Machine Learning era, Special Track on AI for Improving Human-Well Being. The distribution of IJCAI papers in 2019 by area, submitted and accepted is presented at Figure 2, some competitions in Figure 3.

2 Achievements and dilemmas

IJCAI is not only a conference, it is an annual presentation of the world's AI best and brightest and most relevant events, e.g. the meeting of world AI societies. Unfortunately, there are glitches, and this year the presence of national AI representatives was more than sporadic. Hopefully in 2020 the organizer will send an invitation to all national organizations in time. Another idea - as this is an informal organization, we will draw up a list of all national AI societies and submit an invitation in time. Better two invitations than none. In the world of over-aggressive web and other media advertising and social media, activities of scientific societies are somehow overshadowed. For example, while AI funds around the world are growing rapidly with an EU annual increase of around 50%, European AI societies (EurAI) have not increased their memberships, and no new European AI society joined in 2019. When asked why AI societies are not more aggressive in trying to attract more societies and members, the reply was that this can hardly be expected from a scientific organization - e.g. to use commercial advertising methods. Maybe, or maybe not and what is needed are strong vision and determination.

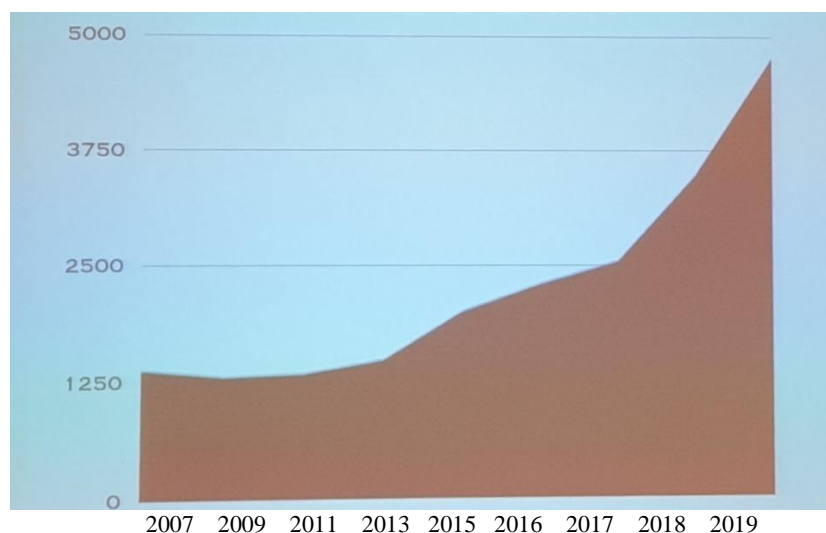


Figure 1: Submissions to IJCAIs.

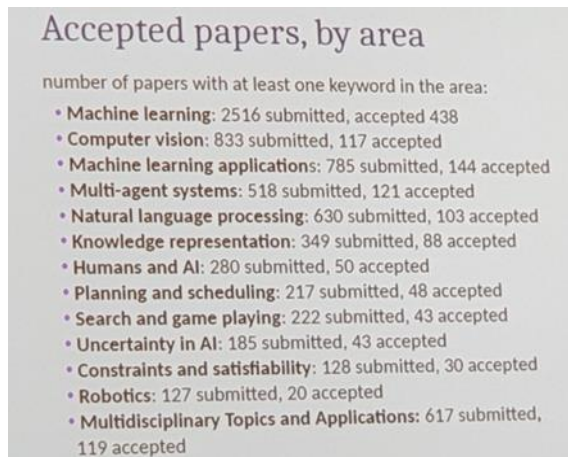


Figure 2: IJCAI 2019 submitted and accepted papers by area.

While IJCAI 2019 as a whole was undoubtedly a great success, some issues were observed in specific areas. For example, "AI in industry" presented several important AI applications. To name one, Xiaowei, a Chinese mobile platform with several AI modules, has around one billion users (Huaawei around 3 billion) and is constantly introducing new AI services. At the 2019 IJCAI conference they presented an AI assistant. When fully implemented and reaching all users, it will bypass Siri (0,5 billion), Google Home (0,4) and Alexa (0,1) users as current leading voice-controlled AI assistants. In the AI demonstration section, a fast growing and flourishing area of IJCAI, several systems deserved and often got world-wide appreciations. For example, one system improved distribution and placement of Uber drivers, enhancing the individual driver's gain and decreasing the user's waiting time. One rewarded demo presented an automatic creation of assistants from websites and the other fair use of workforce.

The industry AI award was given to a Microsoft team for an application of reinforcement learning to personalize news (28% increase in adaptation) and games (40%). These huge improvements were obtained by relatively small modifications of the previous systems, a lesson to be remembered. But surprisingly, the participation in the lecture hall was less than average. The explanation at hand

is that most of the conference attendees were researchers listening to academic presentations that took place simultaneously in several parallel sections. The gap between academia and industry was highlighted once again. Researchers receive funding and fame according to academic criteria and it is not of great importance if their ideas find ways to actually help people and increase profits.

The fusion of real-live applications and academia at IJCAI was courageously attempted in many respects, such as the competition for care of the elderly instead of robot soccer (Figure 4).

The two major advances in 2019 compared to 2018 were probably the increase in massive AI applications, and secondly, new research orientations. The fact that the former was somehow accepted as an obvious fact is not very helpful for AI growth and fame, where an AI program beating humans in a particular game obtains overall attention, while major AI applications hardly ever. But it is precisely the dozens of industry presentations, demos, workshops, competitions such as the elderly-care competition (Figure 4) and practical AI presentations, often related to a particular branch such as robots, that have most impressed an impartial AI observer in 2019. AI is in the intense phase of transforming human society into an advanced, incremental, optimized and multi-objective civilization providing better foundations for long-term sustainable growth.

As usual, there were hundreds of incremental algorithm improvements, be it random forest, boosting or deep neural networks. In particular, the deep learning, where a random forest algorithm is placed instead of a neuron in a network, had shown quite important improvements. The problem, which is consistent with the principle of multiple knowledge [2], is that it quickly loses diversity with additional levels since the random forest consists of more or less all possible decision trees. Therefore, despite some interesting results the original idea of adding another algorithm such as RF instead of a neuron is still in progress. Overall, this incremental progress is quite impressive as AI is used for hundreds of trillions of decisions a day, and a few percent better decisions mean a lot in real life.



Figure 3: Competitions at IJCAI 2019.

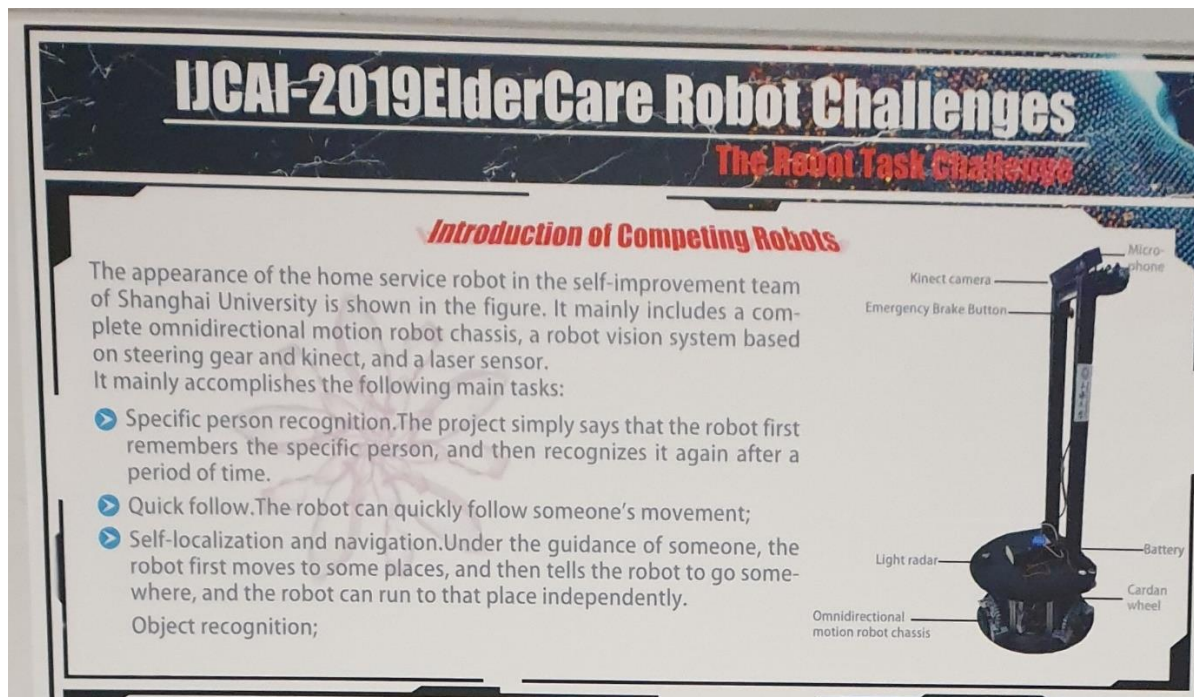


Figure 4: Elderly-care competition instructions. The focus change from academic to real-life was noticeable at IJCAI 2019.

3 AI directions

As far as the new research orientation is concerned, at IJCAI 2019 it was carried out in a rather precise way, although not by all the invited speakers. But several lectures merged into a new research paradigm that was best presented at the conference last afternoon by Broeck, Domingos, and Shoham. It is not a switch from Decision Tress (DTs) to Deep Neural Networks (DNNs) because DTs (or decision graphs) are as accurate as DNNs in some areas and also provide explanations and understanding that are unprecedented for DNNs, regardless of the DNN advances outperforming any artificial system in several areas and often the best people. Nor is it very likely that DNNs will soon become truly intelligent, as Figure 5 shows - a boy standing on a chair hoping in vain to see the stars better with a telescope. It's also not a dilemma whether to use AI systems like GPT-2, HAIM, Grover or "Not Jordan Peterson" (the last one removed due to lawsuits) - why not use them for fun and get acquainted with the power of SOTA AI.

AI is a technology and like any technology it can be used for good or bad. At present, AI it is the one that contributes most to human progress, with applications ranging from robotics to web services and autonomous cars. According to practical statistics, a Tesla car, for example, is nine times safer in the autonomous driving mode than an average classical car. The progress can also be seen in the services openly available on the net - for example, a few years ago nobody could create a system like Not Jordan Peterson - fluently speaking input text that does not differ from the speech of original author.

The actual question / dilemma, according to the IJCAI presenters is the following: Are we on the path to

developing truly intelligent systems or only AI applications capable of playing excellent chess, for example, where the specific algorithmic solutions are dedicated and successful only in a certain area, without explanation and without the impression that something inside resembles a real human intelligence? Since the attempts to solve the Turing test remain as unsuccessful as ever while the computer and AI progress is continuing with the exponential speed, something soberer seems to be hidden in this perplexity.

However, there are the good old strong AIers who claim that we are on the way to true intelligence and we just need to be patient a little more. And anyway, who says that intelligent AI systems need to have human-like intelligence to perform well, because airplanes fly in contrast to birds, while the direct applications of bird-like flight patterns is counterproductive. Why should an autonomous car write or understand sophisticated poems about ethics, mortality or love? Furthermore – consider the moral dilemma of autonomous cars: Who should a car hit - a child or a grandma, if it cannot stop in time? In practice, this is statistically an irrelevant question, since such a situation practically never occurs in an individual's life. Second, in some countries, such as Germany, there is a law that forbids taking preferences based on age - a car that prefers to run over grandma would therefore be illegal and subject to legal consequences. And thirdly, why is this dilemma imposed on the scientific engineering community, where 99% of activities are aimed at developing technical solutions that enable high-quality driving in all possible real circumstances from weather conditions to the reactions of other road users? Should we concentrate our energy on a hypothetical situation or rather design better systems to save thousands of lives




Figure 5: Probably, current deep neural networks, alike other existing AI approaches, will never achieve true intelligence on their own. New strategies are needed.

each year? It should also be noted that two decades ago autonomous cars were more of a joke than something people believed in - so amazing is the AI progress!

So why is a large part of all discussions dedicated to a philosophical dilemma of whom should a car strike first in more or less artificial situations? Why does a Tesla incident caused by an AI mistake attract a lot of media attention and the corresponding ten situations in which it avoided a crash with superhuman reflexes practically none at all? Not to mention that there are several benchmark domains where AI progress can be demonstrated explicitly in a scientific, repeatable and measurable way.

In reality, for a large part of the AI applications, no real human-level intelligence is required and the engineering AI already offers significant improvements. For example, DNNs in the ImageNet benchmark visual recognition test improved their accuracy from 71% to 93% from 2011 to 2014 and from 2014 to 2018 to 98%. To name some of the most important AI achievements in recent years: 2017 - skin cancer, poker; 2018 - SQUAD1.0, Chinese-English translation, Dota2, prostate cancer; 2019 - SQUAD2.0, Starcraft. DNN combined with reinforcement learning enabled a big step ahead. For example, Google's DeepMind played magnificently 50

Atari games. An example game would be hitting the ball from the bottom of the screen upwards to hit objects in several rows at the top of the screen to score points. But to demonstrate the strange nature of some recent AI achievements, when the racket was moved up a few pixels, the performance deteriorated significantly, which is highly unlikely for humans. Another example - when it was investigated how DNNs learned to distinguish cats from dogs, it turned out that their decisions were based on a small number of pixels and when these pixels were changed, with humans still clearly distinguishing between different animals, DNNs failed. These days, researchers are developing algorithms that compete in the search for the minimum number of pixels that need to be changed to mislead DNNs. Several experiments of this kind showed that DNNs learn significantly different than people do, that their knowledge is not general, but highly specialized and therefore brittle. The expectation that DNNs, like other existing AI technologies, will achieve true intelligence is rather an utopianism shown in Figure 5, according to several distinguished AI researchers.

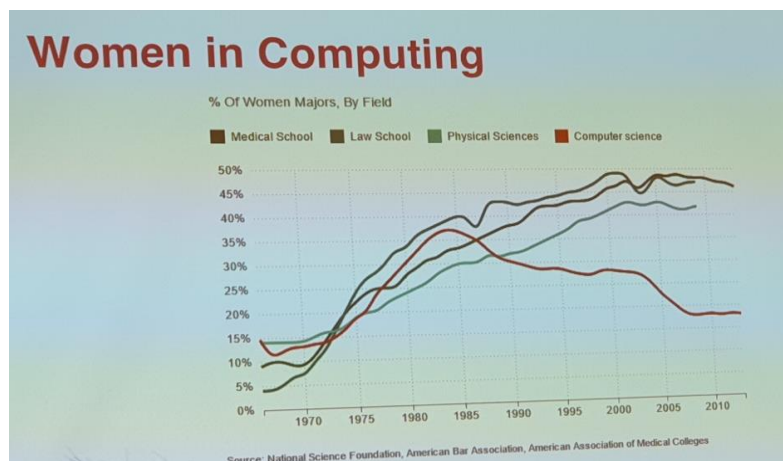


Figure 6: The issue of decreasing number of women in computing is a relevant and sound one, but the idea to use other criteria instead of research excellence for scientific publications is a threat to science.

members with specific personal characteristics, which

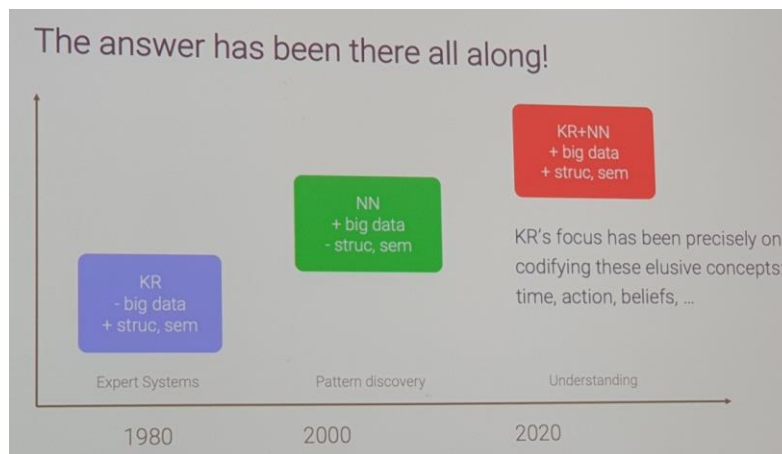


Figure 7: New strategic AI research direction on a path towards true intelligence, proposed by Shoham.

4 AI dangers

Among the dangers for AI, another one lurks in the shadows - the penetration of ideology, which is currently happening in all spheres of human activity, be it mass media or science. Indeed, in some AI areas like superintelligence [3,4], scientific objectivity, i.e. novelty, should not be the only criterion. A hostile superintelligence could harm human civilization, and therefore such attempts should be treated with appropriate supervision and caution[5]. For AI conferences and journals, the standard penetration of ideology is not a major problem, as all these activities are based on anonymous refereeing where authors based on personal characteristics such as position, country of origin, wealth, skin color, gender, age or similar cannot be preferred. In IJCAI 2019, however, there were some attempts to modify this standard approach. They all demanded a "fair" share of a certain section of the population - note that this paper will strictly avoid naming such criteria. Science and ideology do not walk along well, and it is disturbing for any true scientist to observe the growth of ideology in recent years. It might be too early to warn that virtually all civilizations have saturated and collapsed with the growth of harmful ideologies - harmful in the way they collided with the production of vital goods. In Western civilization, the negative effects of overwhelming neo-liberal globalization are quite obvious, from the overburdening of the planet by transporting industries to less developed countries with cheap labor and thus overloading the global environment, to the structure of important positions based on political orientations and personal categories such as gender or skin color, and not primarily on the ability to work well.

Be what it may, the declining number of women in IT technology (Figure 6) is indeed problematic. As a rule of thumb, at least 20% to 30% of the members of the opposite sex are needed to achieve good group performance and according to statistics this is already hard to achieve in some teams. In comparison - for boards of directors, CEOs or ministries, many Western countries require 40% of

represents another extreme, and would be indeed unfortunate if such criteria were introduced into science, as softly advocated. But as long as the refereeing remains anonymous, there is no direct way for ideology to corrupt science too much.

At IJCAI, there wasn't much talk about media IT giants like Google or Facebook, but the way they affect human society to become more polarized, hostile and less open is becoming more and more evident. As an indicative example, suppose one observes a YouTube clip claiming that the Earth is flat. The recommendation system observes the area of interest and recommends more videos that the Earth is flat. Soon all one person will get are videos confirming the wrong belief. There are other stray effects related to the IT giants. Perhaps we should not discuss the penetration of ideology into say Google by deliberately providing the objective algorithms with false data in order to learn to eliminate politically undesirable persons from the media light (e.g. name modifications) as it is probably not a major phenomenon. However, the IT giants with all their positive novelties have also the sinister monopolistic side, influencing elections, and making people more stupid, as the Flynn effect shows with centuries of progress and a decline in the last decade.

At the same time, when used properly, fair and without ideological twists [5], AI will continue significantly improving the quality of the Web. For example, violent videos are removed from YouTube much faster and more efficiently than before with YouTube AI guards. From time to time there are still some failures, such as the elimination of robotic wars with the argument that cruelty to animals exists, but overall the improvement is indeed significant.

5 Conclusion and discussion

To summarize:

- Even without true AI, the incremental AI progress with rather engineering solutions already offers great improvements, and all attempts to discredit them or to shift the discussion outside the scientific, engineering

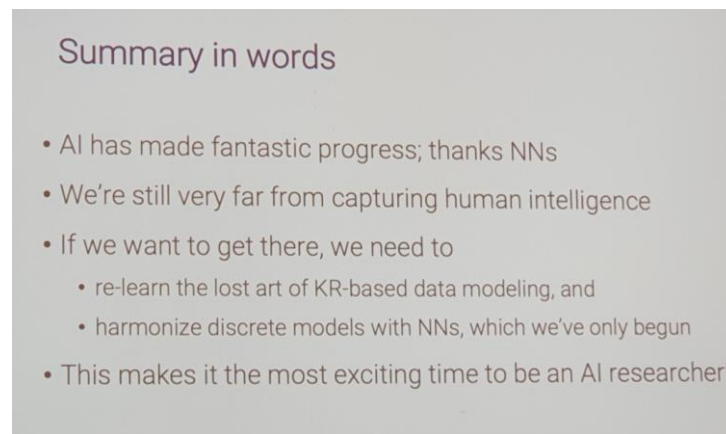


Figure 8: The last page of the IJCAI conference.

solutions are doomed to fail due to penetration of AI into our everyday lives.

- Different authors propose different recipes to achieve true intelligence, e.g., the author of this paper multiple computing and multiple knowledge [2]. In 2019, the IJCAI community proposed to merge deep ML technologies such as DNNs with models - Knowledge Representation (see Figure 7). It is therefore not primarily superintelligence or general intelligence, but the combination of DNNs with model-based (or rule-based) reasoning, that regardless of the upcoming problems at least in the near future will remain one of the dominant technologies and will persist as one of the most important technologies for the progress of human civilization. Believe it or not, even Google's responses to inquiries in recent years have been based on various AI methods. If AI stagnation should occur, then in the form of slower progress is expected, not of the winter type.
- On the other hand, there is no doubt that in a few years' time there will be another financial crisis, because on average crises occur in seven years, and in the last nine years we have been living in a series of steady growth. For the rich and clever, the crisis is an opportunity to enrich themselves more quickly, as history shows, but for the rest of the population, especially the middle and lower classes, there is nothing to cheer about. When and how deep the financial crisis will be - only time will tell.

For the finale of the IJCAI 2019 report and for the overall impression, the last sentence of the Shoham's lecture (Figure 8) was chosen: "**This makes it the most exciting time to be an AI researcher.**"

References

- [1] IJCAI 2019, <https://ijcai19.org/>
- [2] Gams, M. 2001. *Weak intelligence: through the principle and paradox of multiple knowledge*. Nova Science.
- [3] Bostrom, N. 2014. *Superintelligence – Paths, Dangers, Strategies*. Oxford University Press, Oxford, UK.
- [4] Yampolskiy, R.V. 2016. *Artificial Superintelligence*. CRC Press.
- [5] Asilomar principles. 2017, (<https://futureoflife.org/2017/01/17/principled-ai-discussion-asilomar/>).