IJCAI 2015 – The Worst Best Ever

Matjaz Gams,
Jožef Stefan Institute, Jamova 39, 1000 Ljubljana, Slovenia
E-mail: matjaz.gams@ijs.si

Keywords: conference report

Received: August 5, 2015

Editorial

In 2015 the International Joint Conference on Artificial Intelligence (IJCAI) was held in Argentina from July 25th to July 31st. As a result of all the AI hype in the media – and in particular about the exaggerated danger of it ending the human race – the IJCAI was destined to succeed. But was this really so?

With around 30% more submitted papers than at the previous IJCAI in China, and with an absolute record in terms of the number of papers, the conference was a true success. Also, the quality of the papers and the invited lecturers remained superior to all other AI conferences, including rivals like the American AAAI or the European ECCAI. The technical papers showed progress in their particular technical fields, while the invited papers provided a broad overview of the field and the major achievements. Hardly any of the participants would object to these claims.

Interestingly, the number of accepted papers from China slightly surpassed the figure for the USA, and was proclaimed as the no. 1 country in the world - obviously (the EU not being regarded as a single country). Otherwise, the EU would be seen as the leading AI region in the world. One should also keep in mind that China is becoming the world's largest economy, while the EU is a close third. But here it is important to remember that China has a four-times larger population than the USA, and nearly three times that of the EU.

The progress and achievements in AI are not only evident to AI professionals, i.e., researchers and developers, but also to the public in several ways. For example, more and more companies are intensively investing in AI and, besides hiring AI staff, also provide substantial incomes. There is hardly any major IT or web-related company that is not using several AI methods. In practically all areas of computer and mobile use, AI is introducing new functionalities.

To demonstrate this concept, one could start with music, as was performed live at the opening of the IJCAI. Anybody from the audience could select a note in the musical scores, a pianist reproduced the selected music from that note on, and within a second or two the computer recognized the composition, displayed the current bar and began following the performance in real-time. Fascinating! Congratulations to MusicCompanion, the Austrian Gerhard Widmer and his team. This example nicely indicates the new orientations and achievements of modern AI, i.e., everybody is familiar with the possibility of computers finding any specific text on a disk or in a book, large directories or large databases. To achieve the performance of MusicCompanion, the sounds of a piano, or any other form of audio, have to be transformed into a computer-readable form, and then a proper program can find its location in the databases of properly represented music events – from text search to music search in just a couple of years. There was no major breakthrough in computer or artificial intelligence, "only" very clever programming, good AI and several years of hard work. Similarly, the list of such successful applications applied to real-world problems goes on and on.

The Computer and Thought Award recipient Ariel Procaccia presented his achievements in three areas. The first was in kidney exchange (human-organ transplantation), where AI enables a significantly better exchange due to its mechanisms: intelligent fight with combinatorial explosion, successful introduction of time and the ability to rapidly incorporate modifications. This and other achievements have been widely published in the USA and the world's media. A similar case exists with security scenarios, where the task is to organize patrols in such a way as to optimize the defence against attacks. Millind Tambe is recognized as the pioneer in this area. With modifications, these AI methods are being deployed at various national institutions. The third area is fair division, e.g., of a heritage. A computer program finds the best solutions given user preferences. The optimality is not guaranteed theoretically, yet it is very often achieved in reality. This and several other AI programs are available through the web.

The list of other prize recipients includes Barbara Grosz (Research Excellence in NLP and Multi-Agent Cooperation), Bart Selman (John McCarthy Award for Taming Complexity in General Inference Mechanisms) and Anthony G. Cohn (Donald E. Walker Distinguished Service Award for Contributions in AI Societies and Journals).

There were so many interesting papers and presentations that any IJCAI scientific reporter has to rely on some personal bias. For example, Christof Koch presented Consciousness in Biological and Artificial Brains. First, he described the necessary conditions for (human-type) consciousness, and then based on reported human cases provided proofs that one needs no feelings, no self-consciousness or several parts of the brain to be conscious. The minimum that is needed is approximately one hemisphere of the human brain, with the

corresponding neocortex of a live human. For example, the cerebellum, sensing or other brain or neuronal parts were found to be missing, but with the human still being conscious. From this minimal functionality Koch presented five axioms and the related consequences, and finally proposed a formula for computing the amount of consciousness, producing a numerical value as the output. According to his presentation, the formula for consciousness relies on the number of nodes and the complexity of the neural network's architecture with recursive loops and modules. As a consequence, current computers probably do not enable a physical realisation of consciousness. In his opinion, current computers therefore have zero consciousness, while animals have some and humans are at the top of the list. Future computers will sooner or later achieve consciousness, but the architecture will have to be physically different, not just different in terms of simulation. In the same way as a computer simulation of a black hole does not bend the time-space, a computer simulation of consciousness cannot produce true consciousness. These concepts might be hard to grasp, but just consider that a prototype of a car moves in a real world, while a simulation of a car does not move in the real world. Similar relations are valid for computers, the mind and consciousness. In any case, these revelations provide several new possibilities for thought and research experiments in AI and the cognitive sciences.

The panel "Rethinking the Turing Test" was based on the need for another, better test of computer intelligence. The majority of speakers revealed weaknesses of the Turing Test (TT): the test is about human psychology, TT competitions are about tricks designed to deceive an average interrogator with mistakes, deviations, etc., not helping to develop new AI methods. To measure actual AI progress, another test is needed, e.g., the Winograd schema, with current scores of about 70 (they did not mention that for random replies, the score should be 50). There are other tests, like Captcha, that can distinguish software agents from humans on the web because the software is not able to read blurred text. The Lovelace test is based on creativity capabilities. Another test deals with impossible stories or pictures the computers have problems with. All the panellists agreed that improvements beyond the TT are needed. The most damaging opinion was provided by a Russian child, who asked "What would our world look like without the TT?"

Maybe the only interesting confirmation of the TT was that no program has ever passed the Turing Test and that the media reports about such an event in 2014 were not well founded: "The 65-year-old iconic Turing Test was passed for the very first time by the computer programme Eugene Goostman during a Turing Test in 2014, held at the renowned Royal Society in London on Saturday. 'Eugene' simulates a 13-year-old boy and was developed in Saint Petersburg, Russia." These statements were later dismantled in a scientific paper and other publications. The trick was in misguiding the interrogators by telling them that the computer was a 13-year-old boy from the Ukraine in order to explain the odd responses. Looking at Eugene's replies it seems quite naïve that any human fell for it, but then again looking at situations when I was pickpocketed it seemed quite silly as well. It is important to realise that neither of the two situations had much in common with the true nature of the TT.

It seems that Stephen Muggleton was the only person objecting to new tests: "Cheetahs run faster than humans, but are not more intelligent. Similarly, these new tests will not measure intelligence."

My opinion is that there is nothing wrong with proposing new tests; however, there is no need to misinterpret the TT, one of the major scientific tests. The true concept of the TT is far beyond the understanding of the mass media. For example, consider that the students during one of my classes are instructed to propose a couple of questions that would reveal the true identity of a human/computer. Even though they are warned at the lectures that these questions will be asked at the exam, the students find these issues hard indeed. They typically prepare several questions in advance, but stall when the task is slightly modified or directed to enable perceiving a true understanding.

The catch is that computers have zero true understanding, semantics, and zero top human mental performance. If a new statement is formed and the next question is about the meaning of the previous sentence, computers fail miserably. The TT therefore correctly shows that current computers possess no true intelligence, as does Koch's equation for consciousness. Any test showing current consciousness or other top human mental property would clearly be bogus. Designing such a test would be a matter for witchcraft, not a matter for science.

And the question about what a world would be like without the TT – this is similar to several other scientific comprehensions and laws, e.g., about the existence of black holes or that there are other planets in our universe. Without this knowledge we would not properly understand the world around us. For example, without the TT we might think that we already have human-smart computers and would have no need to design novel AI systems. Then we would be like a civilization designing better and better springs in an attempt to travel the universe, not realising that we need something entirely different.

There have been several areas of significant AI progress recently, among them the smart assistants like Siri or Android Assistant. At the IJCAI, several new systems were presented, for example, for helping the elderly at home alone with a specialized knowledge of certain needs and tasks. For example, a Chinese robot, Vivian, has knowledge about travel, tickets, banking transactions, etc. Facebook presented cognitive computing as the new paradigm, consisting of learning for personalised information, and not forgetting IBM, which was involved in several new AI research activities, including previous victories over humans in chess and Jeopardy, and which produces several patents per day. Companies like Microsoft, and in particular Google, presented their AI achievements. of which several are

well popularized. The "Deep machine learning" term is often related to recurrent neural networks, but has progressed into another generation of machine learning, the one that renders ML repositories like Weka a kind of obsolete. The shift has moved from users using specific algorithms towards user declaring the task and the AI system chooses the best methods and parameters to fulfil the task. And is it not Auro-Weka which searches for optimal setting of the particular algorithm; it is also dealing with the algorithms and parameters and not with the task itself. Another less well known event is the Alibaba success story.

The Chinese company Alibaba was established as a private company with one owner in 1999. Their goal then and now remains the same: making business easy anywhere. Currently, they are the world leader in many parameters, e.g., in the number of registered users, the number of transactions per day, etc. One of their products is Taobao, the 20th most often visited webpage with 800 million products available for purchase. Several of their services are based on AI systems, mainly on machine learning. Among them is the dynamic-pricing, credit-scoring system and smart logistics. The credit system is totally autonomous – based on data about a particular customer it proposes a well-augmented amount of credit or rejects it. According to reports, the system will become the new standard credit system in China.

What about the media hype that concerns AI destroying the human race? As could probably be expected, no technical paper and no workshop or tutorial focusing on that issue took place at IJCAI2015. However, a couple of authors mentioned it during their presentations. For example, Ariel Pocacca, the first awarded speaker, rhetorically asked where the media sees a correlation between the AI contributions and the dangers to humanity. He presented three systems that are clearly beneficial to humanity, e.g., the first one about improved kidney exchange. On the other hand, the discussion about autonomous killing machines was one of the central debates in the non-technical presentations.

An open letter, calling for a ban on offensive autonomous weapons, signed by Stuart Russell, Nils Nilsson, Barbara Grosz, Tom Mitchell, Toby Walsh (contact person) and 1000 others was released to the press on the first day of IJCAI 2015.

Russel Norvig in his lecture asked the audience what a good AI researcher should do in a situation when the media reports the potential dangers of AI. While the researchers at the IJCAI present world-class achievements that actually move AI and humanity forward, the media remains ignorant of these true achievements and propagates the ungrounded fears proclaimed by non-AIers. His advice is to relax, as such mortal fears follow any major new technology, while the media propagates the most attractive views without understanding whether there is indeed any cause for concern. As long as the field is progressing, the salaries increase and the number of students grows, we have no reason to complain. AI research is more successful than ever, consistently finding ever better transitions from scientific research into the real world. influencing our everyday lives along the way. In his words: "We are faster and faster moving towards the greatest event in human history". That, of course, is the singularity point, where truly intelligent programs and machines will propel human civilisation into a new era.

At the panel organized by the conference chair Michael Woolridge (Oxford), the invited participants expressed rather diverging opinions. The announced panellist list of Maria Gini (Minnesota), Barbara Grosz (Harvard), Francesca Rossi (Padua), Stuart Russell (Berkeley), Manuela Veloso (CMU) with the agenda "who speaks in the name of AI" was soon distracted by the ban initiative. "We have a bewildering array of different organisations at the national and international levels representing us (AAAI, IJCAI, ECCAI, PRICAI, KR, etc.), with very little coordination or communication between them. Researchers in distinct sub-fields often work in their closed worlds, unaware of the work that is going on in other sub-fields of AI, and the development of the field is hindered by endless fragmentation…, the lack of any authoritative voice for AI creates a vacuum, where ill-informed speculation about the potential of AI is rife, and attention-seeking claims in the popular press receive unwarranted attention, with nobody in a position to speak for the field, and to give an authoritative, informed, and balanced response."

It's perhaps worth mentioning that there was no "authoritative, informed, and balanced response" from the panel as well. One panellist claimed that the fear hype is providing negative implications for AI because people will associate negative feelings with AI. Another panellist claimed that AI should not get involved in any discussions about such mass-media issues or anything that involves politics of any kind, yet, the woman issue was often reintroduced. The initiative by Russel and Toby Walsh to submit a petition to the UN to ban autonomous killing machines, i.e., without a human decision in the loop, was also greeted with several remarks. Yet, overall, the majority was not against the petition, rather the way it was presented. The IJCAI, after all, should be a democratic institution, with the participants as the voters. If a group of individuals stands out that may be because they want to draw attention to themselves – as was mentioned in one of the remarks. Another objection was that a strong scientific discussion is needed before any petition. In summary, while the debate was not convergent, the overall impression was that the majority of researchers favour the ban petition, albeit with more elaborated procedures.

Despite initial problems, the AI community is wakening up. Indeed, it seems strange that other professions play a key role regarding AI-related issues. Probably, the discussions at IJCAI prompted activities in other societies. For example, in Slovenia some scientific societies already supported the ban and voting is going on in several more. In a reasonable time, the Slovenian societies will submit their support for the ban to the national government and also to the UN.

According to the title of this editorial, the worst issues of the best IJCAI ever should also be mentioned. The purpose is clear – by highlighting potential

improvements, the next IJCAI could be even better. For one thing, the initiatives to come to an agreement about how to proceed with the ban petition and the agreement about how to propagate the voices of the AI researchers were at a rather basic level, without an actual procedure being determined. Yes, there is an agreed need, but proper formal procedures are needed as well.

Second, the organization at the hotel was peculiar at best. During the first lecture of the conference, renovation was going on in the next room, with banging and drilling, sometimes rendering the audio presentation incomprehensible. Third, the air conditioning in the rooms was kind of random – in some halls it was as cold as in the ice age and in others there was practically no air conditioning, leading to jungle conditions. That said, the Argentinian local organization proved to be very friendly and supportive.

Worst of all, due to there sometimes being as many as nine events in parallel, it was often impossible to visit the most interesting lectures. Whoever constructed such a program, putting in parallel several of the most relevant presentations, would certainly need to reconsider the scheduling next time. Or maybe an "IJCAI lemon" should be delivered with respect to this issue, in order to prevent a repetition in New York, the venue of the next IJCAI.

For those of you that are not familiar with the so-called "pig-style" events, one should attend the IJCAI cocktail party after the opening ceremony. The term comes from the analogy with pigs messing at a scarce food source, and later eating in a crowd standing and bumping into each other. The definition was fulfilled. One might argue that the same was true at the regular coffee-breaks, but there were no rows in front of the coffee and water stands. However, the lack of any food, even modest cookies, left a miserable impression, especially having in mind the unavailability of the published invited and awarded lectures and the enormous conference fee, which approached €1000 for the workshops, tutorials and the main conference. Surely, it was an enormous effort to organize together so many relevant events, but the room for improvement is quite large.

At the same time there were several improvements compared to the previous conference. For communicating AI matters to the media, the IJCAI ran a daily press conference, livestreamed to the world every day. This was a first for the IJCAI. The proceedings were published on the web two weeks in advance, the program was on schedule, and a large majority of the lectures were comprehensibly presented. The number of contributions is increasing, leading to the singularity at best, or something like that in the worst case. The IJCAI2015 had plenty of lectures to listen to and achievements to admire.

Regarding the representatives from Slovenia (remember, Informatica is also supported by the Slovenian SLAIS), there were just two issues worth mentioning. We co-organized one workshop, and at an ambient-intelligence (AmI) tutorial Juan Augusto described the PhD of Hristijan Gjoreski as one of the major events in ambient intelligence. Gjoreski designed an AmI-related version of a general random forest algorithm, improving the accuracy at statistical levels, at worst, and enormously for some demanding AmI tasks. The reason why Slovenia is lagging behind its usual achievements at the IJCAI is that we still attend the ECCAI because of European relations (to see and communicate with Muggleton, de Raedt, de Mantaras and other key scientists and institutions in the EU), while the national evaluation procedures render all IJCAI achievements (papers and their citations) as practically irrelevant. A couple of years before it was possible to publish a paper in a good journal from the IJCAI conference, but now the conference publication itself prevents similar journal publication, and in the national evaluation system a major achievement gets no acknowledgement. Hopefully, publications like this will help change bureaucratic irregularities.

In a conclusion, while AI is no doubt a current and long-time success story and a future hope for humanity, there is a lot of space for improvements, e.g., by organizing procedures to transmit AI opinion to the media. In all scenarios, staying open-minded, democratic and first of all true scientists and developers, it is up to the AI society to fulfil its own prophecy and not the mass media's fears. It should be a classic case of a self-fulfilling prophecy!