Benelux A.I. Newsletter

AI and Robotics

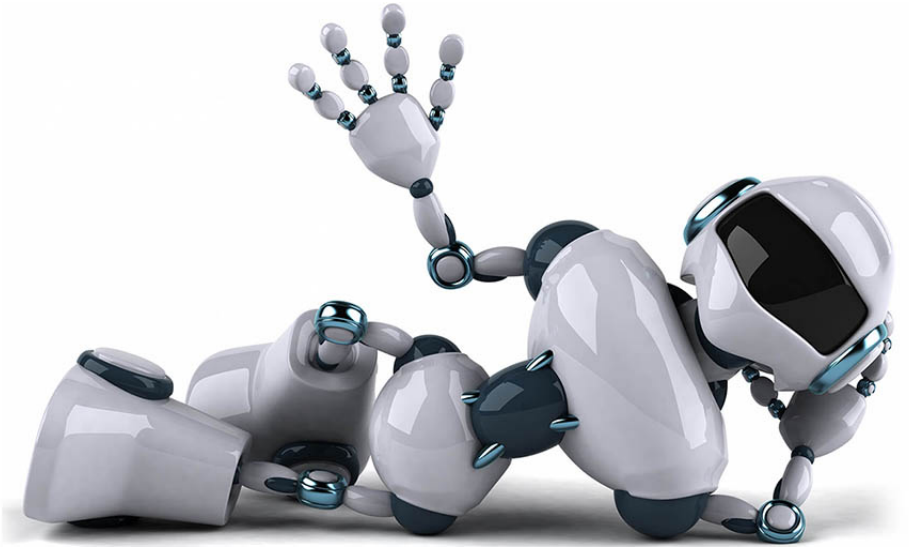

Summer 2016 Edition (No. 2, Vol. 30)

BNVKI
AIABN

Contents

Interview with... Manuela Veloso

by Marc van Zee

Profile. Manuela Veloso is the Herbert A. Simon University Professor in Machine Learning, Computer Science, and Robotics at Carnegie Mellon University. She was the President of AAAI (Association for the Advancement of Artificial Intelligence) until 2014, and the co-founder and a Past President of the RoboCup Federation. She is a fellow of AAAI, IEEE, and AAAS. She is an international expert in artificial intelligence and robotics.

What is your academic background and how were you drawn to the questions you are dealing with now?
I was born in Portugal, where I did my bachelor studies in Electrical Engineering. After that, I moved to the United States in 1984 where I did a master in Computer Science at Boston University. Then I started a PhD at Carnegie Mellon University, which I completed in 1992. I became a faculty member at CMU in 1995 and I stayed there ever since.

Manuela Veloso

This means you have experienced research and university life in both Europe and the US. What are some differences between research in the US and in Europe?
At CMU, and in the US in general, we do mostly basic research. Although we have applications in mind, the advances to the field are fundamental. For instance, at the time when I started with the RoboCup project there was no clear application. The concept of teams and robots was very new and vague, but I was able to do research on it nonetheless. I did not have to show an application, it was just an exciting thing to do. I think that the ability to be free in terms of what you want to pursue is a good thing about research.

Do you currently have any collaborations with universities in Europe?
I am part of a CMU/Portugal program for the last ten years, in which universities in Portugal have collaborations with CMU by jointly supervising students and organizing joint projects. Because of this, I have had a lot of interaction with researchers in Portugal in the last ten years. Moreover, because I was president of RoboCup, I have a lot of contacts with universities all over Europe. We collaborate in AI and robotics.

> “In the coming years, everything that is going on inside a robot (probabilities, decisions, etc.) will have to be translated into language.”
>
> *Manuela Veloso*

What are some researcher who have inspired you in your research?
For me, the work of Herb Simon and Allen Newell on problem solving in the mid 50’s was fascinating. Before that, Computer Science was about very specific areas such as cryptography or algorithms. Newell and Simon defined, in general, how a computer can solve any problem. You have an initial state, a set of available actions, and an objective function. The task of the program is to end up in a state in which the

objective is satisfied. Topics that are currently very popular such as planning, search, learning from experience, and advice all started from there. I still think in this way in my current research.

A prominent part of your current research is about robots asking for help from humans. Can you tell us something more about this?
I am passionate about autonomy; about how robots can make decisions in the real world. In my PhD thesis I worked mostly on planning, but after that I started focusing more in robotics. One thing that we learned about eight years ago about autonomy is that we currently do not have, and perhaps never will, have the technology to develop robots that can do everything alone, because the world is simply too complex. I realized that a logical way to solve this is to develop some kind of symbiosis between robots and humans. In our department we have several collaborative robots, called CoBots. When the CoBots can't do something they proactively ask for help.

Robots that realize that they cannot solve something alone. This sounds a bit like the concept of bounded rationality of Herb Simon. Is this a coincidence?
No, this is not a coincidence. I grew up in the tradition of suboptimality, statisficing, not optimization. Because I had this background, the idea of asking for help came very natural.

So how does a robot know when to ask for help?
This is currently hard-coded into the CoBots. For instance, the robot asks for help if it doesn't have the necessary actuators to carry out a task (e.g., it cannot open a door). But it for instance also asks for help when it doesn't understand what the user is saying. This means that there is no mapping from what you are saying to concepts it has in its knowledge base.

Do the CoBots also use the internet?
Yes, they do. Mehdi Samadi, one of my PhD students, was involved in a project about using the web to find out which pairs of entities occur most often together in order to, for instance, locate coffee.[1] In this way, the robot develops confidence in the truth of propositions. I could for instance ask the robot to find coffee. It can learn from the web that coffee occurs often in the kitchen, so it would go to the kitchen and look for coffee.

Can you tell us a bit about your current research?
I recently became fascinated on how to make robots more introspectable in order to do some diagnosis. It is very hard for a robot to answer questions such as: What just happened? Where did you come from? Why are you late? Who were at the elevator? Of course, all actions and events are logged in a robot, 60 times per second, but it is difficult to extract the relevant information from this. My last work has been on verbalization: The ability for a robot to look at its own data and verbalize the relevant experience. In order to do this, we created a verbalization space with different dimensions such as abstraction, locality, and specificity. You can ask the robot: What happened in the elevator?

It sounds like dialogue / natural language plays a big role in this.
Yes, this is absolutely necessary in order for robots to become trustable! Autonomous robots currently don't exist. Your cellphone is not autonomous, it doesn't disappear from your sight. When my CoBot turns around the corner, I can't see it anymore. I want to know what happened. This is what will make them

[1] The thesis can be found at `http://www.cs.cmu.edu/~mmv` and papers at `http://www.cs.cmu.edu/~mmv/Veloso.html`.

more trustable, accountable, and robust. In the coming years, everything that is going on inside a robot (probabilities, decisions, etc.) will have to be translated into language. So I'm working on accountability, trustability, how to talk back.

You now seem to focus on robots communicating with humans, while before, in the RoboCup project, you focused on robots communicating with each other. Why?
At first, it was difficult enough to get autonomous robots working without human interaction. Now that we have the CoBots working with each other, we started getting interested in interaction with humans. It was a logical next step. They also still work together.

> People will have to switch their minds to enabling robots not being very good in what they do. If people insist on a fully omnipotent robot then we will never have robots around, I am very pessimistic about this.
>
> *Manuela Veloso*

You are head of the machine learning department of CMU. Machine learning is currently very popular in industry. Do you think machine learning is the answer to all of our problems, or do you think there is room for symbolic / logic-based approaches as well?
Machine learning captures all the concepts I was using for a long time, most importantly improving with experience. A lot of machine learning includes statistics, but it also contains logic, extracting relationships, rules, inference from data. We all use logic. My planning operators use logic. We use logic in (probabilistic) rules. We all represent knowledge in one way or another, from bayesian networks, to dependencies networks, to neural networks. They are all simply representations, and some researchers are more drawn to specific fields than others.

How do you see the future of robotics? When will we see them in our every day lives?
There is no reason for us not have robots like CoBots in our daily environment. People will have to switch their minds to enabling robots not being very good in what they do. When you buy a car or a phone, you expect it to work flawlessly. All our technology: refrigerators, toasters, computer, they all work. We have to learn to buy a robot that doesn't work well in the beginning. The AI system should be about becoming better. You buy something that does not know what you want because you have to train it. This is necessary, because a robot cannot know everything, there is too much!

Then maybe we shouldn't start with risky applications like self-driving cars?
Yes, we should have something in between the Roomba and the self-driving car. I believe that is it essential that we have moving things. I love Amazon Echo, but it doesn't move. Your cell phone doesn't move. There are very few moving things around in our daily life.

What's Hot in... Ethics and Philosophy?

by Robert Jan Sips and Zoltan Szlavik. IBM Benelux Center for Advanced Studies

AI, and in particular, learning systems have been receiving increasing interest in the media and among the general public as well. It has been already five years since the IBM Watson system won the Jeopardy challenge, and attention to the field has been growing since. Google DeepMind's AlphaGo, systems that can chat, 'dream', generate speeches, or create paintings are easy to relate to, and so, more and more people have been talking about what AI, specifically, learning (aka. 'Cognitive') systems, can do.

As computational power, as well as code (see the abundance of it on github), is becoming more easily available, we can all experiment with the new technology, train and use 'pre-cooked' deep neural networks, and see what we can do with them. We can have fun: Do you want the computer to generate a unique name for your to-be-born baby? Here it is, within an hour, based on Dutch first names in circulation: Donty or Roske if it's a girl, Teike or Devino if it's a boy. Would you like to be inspired by the wisdom of the machine? How about a quote by 'Necessious' (as it named itself): "Intelligence is great darkness."

In the corporate world, too, we have been moving to more easily accessible and usable, cloud based services. Many parts of the Jeopardy Watson system are now available through cloud APIs, and new services around Cognitive Computing are being added to that portfolio regularly. Combining these services, we can build clever systems, and build them reasonably quickly and easily.

The above examples illustrate how we can do a lot easily with the new technology, and one may say that Artificial Intelligence is commoditising. However, with this wide availability and low entry barrier, it is easily forgotten (or misunderstood) that technology has boundaries, both ethically and technically. The latter is the easier one of the two. If you use the technology in a way that does not make sense, it will simply not work well enough, and the potential business that you might want to build around the idea will probably meet the harsh reality of no incoming funds.

However, what if an idea works -and the abundance of computing power and available data makes more and more ideas viable- , but either it is for a questionable purpose, or it works but not equally well for all who would be your customers? We are often forgetting about the ethical (and sometimes legal) considerations of using the new technology.

What if something goes wrong? Who is responsible, what is our role in dealing with the consequences? Can we, and should we use data just because we can access it, and our system might get better classification accuracy if we use it as training data? How can we prevent the wrong kinds of mistakes, and avoid biases that are introduced into Cognitive Systems? There are more and more examples that can serve as excellent learning opportunities in the context of these questions: Microsoft's Tay going bad is such, the demo of Volvo's self-stopping car is a similar one, DeepMind's using of data from the National Health Service is another one to talk about. Note though that it could be any organisation, large or small, academic or industrial, that can produce such easily and widely visible examples. As long as we learn from these and avoid them in the future, they were perhaps mistakes worth making.

The above mentioned questions are often very difficult to answer, but, thankfully, more and more organisations are beginning to realise that they need to address them. For instance, to use a local example, the Dutch Authority for Financial Markets (AFM) has these questions on their agenda and is seeking advice on the boundaries of "Robo advice". They would like to know what learning systems mean for the financial world, and how to deal with the ever more rapidly developing technology that is increasingly hard to follow in terms of policy making and legislation.

Likewise, the Driving Vehicles and Driver authorities (RDW and CBS) have started discussions (with us amongst others) on how to deal with the onset of self-driving vehicles and how to deal with them.

Where are most of the mistakes made, and where do unwanted biases come from? Where do we need to be careful?

For AI scholars, these questions and many of the potential answers are naturally not novel. We know that learning systems make mistakes by design, that data and the annotations thereof induce bias into the foundation of these systems and that the outcomes or advices based on these systems should be considered with this knowledge in the back of our minds.

Moreover, the debate on the foundation of intelligence and the philosophy around it (Can a system have conscience and make conscious decisions? Can machine intelligence exist?) is fundamental to the role these systems may play in society. This debate is more alive than ever. During the BNAIC 2012, we were present at a keynote given by Chris Welty (IBM Research, now Google), on the Watson Jeopardy challenge. Jaap van den Herik confronted him with the question whether a system like this could be considered intelligent. The answer was (and this is one which is gaining momentum in the day-to-day talk in industry) that the machine made very stupid mistakes, when observed from the perspective of a human. Likewise, if we were able to ask the machine, it would say the same about the human. Machines and humans think differently and exhibit a different "intelligence". For instance, 'Necessious' has no understanding that it is writing "quotes", likewise, Google Deep Dream does not "know" it is generating art-like images. There is no creative inspiration; yet, if people pay thousands of dollars for art, like they did in a recent auction for art via one of Google's Deep Dream networks, who are we to judge the quality of the "mind" behind the painting, or why should the general public even care about this judgment?

We would like to use this medium to draw your attention to two initiatives around ethics and awareness which have surrounded the ECAI conference in The Hague: the first is the AIckathon, which will host the "IBM Challenge": how can we create an automated system that explains AI and the ethical issues surrounding it to the general public, using past ECAI proceedings. The second is the IEEE working group on Responsible AI, mostly consisting of corporate researchers (IBM, Cisco, Google) which has presented its report during ECAI this year.

Despite warnings about a potential next AI winter, we do believe there is a healthy future for AI, if used prudently and ethically, for the right purposes. Open collaboration and transparency is the way forward. Or, to put it in the words of Necessious: "A happy peace is my favorite vision".

Know your Classics: Fuzzy Sets and Fuzzy Control

by Ann Nowe (University of Brussels)

In this first edition of "Know your classics" I would like to take you back to the original publications by Lotfi Zadeh (born in 1921) on fuzzy sets and related concepts. Fuzzy sets [L.A. Zadeh, 1965] extend the characteristic function of classical sets, which maps elements into the set $\{0, 1\}$, to the interval $[0, 1]$[2]. This basically means that an element no longer fully belongs to a set or does not at all belong to the set. Fuzzy sets allow an element to be a member of a set to a certain degree. If we consider the set of the tall people, in a crisp or classical sense a person of say 1.70 meter would not be considered to be tall and therefore this person would not belong to the set of tall people. In a fuzzy set context this person might be considered to be part of the set with a membership grade of say 0.2. A person of 2.10 meter would be fully agreeing with the label of being tall and therefore be assigned a membership degree of 1.

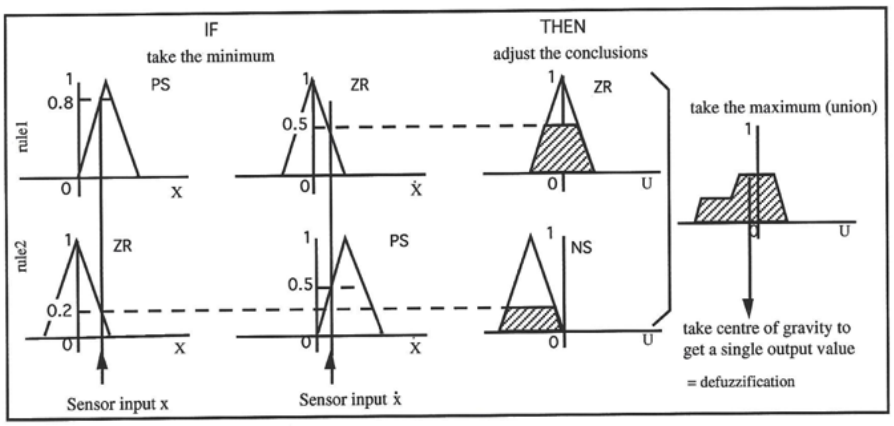

Figure 1: (View in browser) Graphical illustration of the inference in fuzzy control. The membership grade of the inputs (here x and $\dot{x}$) with respect to the fuzzy antecedents are determined (resulting in respectively 0.8 and 0.5 for rule 1 and 0.2 and 0.5 for rule 2). The minimum operator is used as and operator to combine the membership grades for the different antecedents (resulting in an "activation" of rule 1 of 0.5 and rule 2 of 0.2. In this example, the minimum operator is also used as implication operator, resulting in a cut off of the fuzzy conclusions and respectively 0.5 and 0.2, shown by the shaded area in the fuzzy conclusions. The resulting fuzzy conclusions are then combined using the max operator, shown at the right hand side of the picture. To finally obtain a single control action that can be applied to the system defuzzification is applied, here the center of gravity is applied.

Fuzzy sets, and in particular fuzzy controllers which are based on fuzzy set theory, were a real hype in the late eighties / early nineties. Some people loved it, some thought it was just something passing by. To this latter group belonged, Charles Elkan, who published the award winning paper at AAAI 93, "The paradoxal success of Fuzzy Logic"[3] In this paper he "proves" that fuzzy logic reduces to binary logic. However, it was common knowledge that within a fuzzy set framework, and by extension also fuzzy logic, one always has to give up at least one of the laws of Boolean logic, such as for instance the law of the

[2]More complex variants exist but will not be mentioned here.

[3]`https://www.aaai.org/Papers/AAAI/1993/AAAI93-104.pdf`

excluded middle. And it was exactly this law Elkan used implicitly in his proof. In the same paper, Elkan recognizes the success of fuzzy control, because of it's simplicity. However, to my opinion, it is fuzzy control that has led to justified criticism against fuzzy logic. Despite of the very many patents that were filed based on this theory, witnessing its industrial success, there are some issues concerning fuzzy control I want to share with you.

Fuzzy control is based on a set of fuzzy if-then rules[4], such as *if the state of the system is medium positive and the velocity is small positive then the control action is medium positive.* The *and* operator and the *implication* as well as the combination of rules are performed using fuzzy set operators. Depending on the exact choice of the operators being used (there are complete families of possible operators that possess different properties) the concluded control action corresponding to a specific state measurement slightly differs. The most commonly used combinations however result in a graphical interpretation a shown in Figure 1. If we look a bit closer to the mapping these rules express, we see that the width of the fuzzy conclusion is a second class aspect. Meaning that the impact of the width of the fuzzy conclusions is either complete obsolete or has a very minimal impact. This is illustrated in Figure 2. This figure shows a typical single input single output (SISO) case. If input is *Negative small* then output is *Negative small*, if input is *Zero* then output is *Zero* and if input is *Positive small* then output is *Positive small*. In green we show the control law that results from this rule base if the *product* operator is used as *and* and *implication* operator is used and the *bounded sum* to combine the different conclusions. If the operators *min, min* and *maximum* respectively are used then we observe a very similar behavior. However not quite linear, but the "global" behavior is clearly linear[5]. If we change the width of the fuzzy conclusions then in the "green" case there will be no impact at all, in the "red" case then we observe slightly bigger local deviations from the linear behavior (red dashed line). In this example the behavior is globally linear, but by shifting the fuzzy conclusions on the y-axis the global behavior might become piecewise linear.

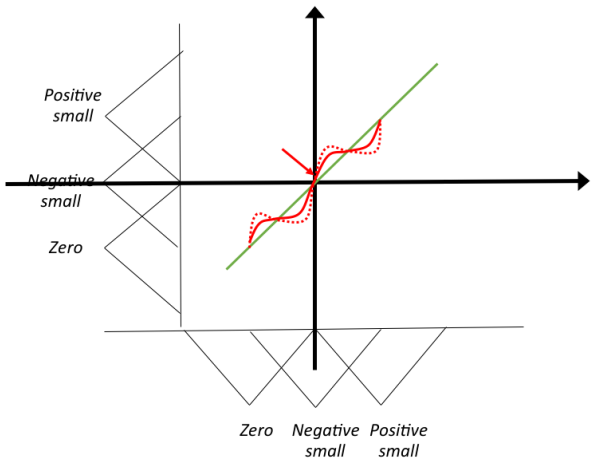

Figure 2: (View in browser) The green line shows the resulting control law when the product operator is used as and and implication operator is used and the bounded sum to combine the different conclusions. The red line shows the resulting control law when the operators min, min and maximum are used resp.

[4]We restrict ourselves to the Mamdani type of fuzzy control.

[5]More details are provide in Synthesis of "safe" fuzzy controllers based on reinforcement learning, PhD Vrije Universiteit Brussel, Ann Nowe, 1994.

Given the fact that the width of the fuzzy conclusions does not seem to have a big impact, one can justly wonder what the advantage is of using fuzzy conclusions. The reason why fuzzy controllers in some settings were shown to perform better that their PID counter parts, which are also linear controllers, is due to the small non-linearity one observes in the case the combination of *min, min* and *maximum* operators are used (see red arrow pointing at the non-linearity in the red line in Figure 2). This small extra force in the control action makes the system just that little bit more stable.

A quite different interpretation of the rule base of a fuzzy controller, is to view the rules as expressing a *fuzzy graph*. Doing so we interpret the rule if A then B, with both A and B fuzzy sets, as "A is coupled with B" rather than "A entails B". This entirely agrees with a remark made by Zadeh on the original fuzzy implication. In [L.A. Zadeh, 1975] the fuzzy rules of the form if A then B are treated as a special case of if A then B else C. The former rule results from the latter by assigning to C the universe of discourse. It is remarked that another option could have been to set C to the empty set, which reduces if A then B else C (defined as $(A \times B) \cup (\neg A \times C)$) to the Cartesian product $A \times B$. Doing so B expresses a fuzzy restriction on the allowed control actions if the state belongs to the fuzzy region described by the set A. A more narrow fuzzy set B expresses that the control in the region A is more critical, while a more wide set B expresses that there are more control actions acceptable for the region A. Figure 3 shows a graphical representation of a fuzzy graph.

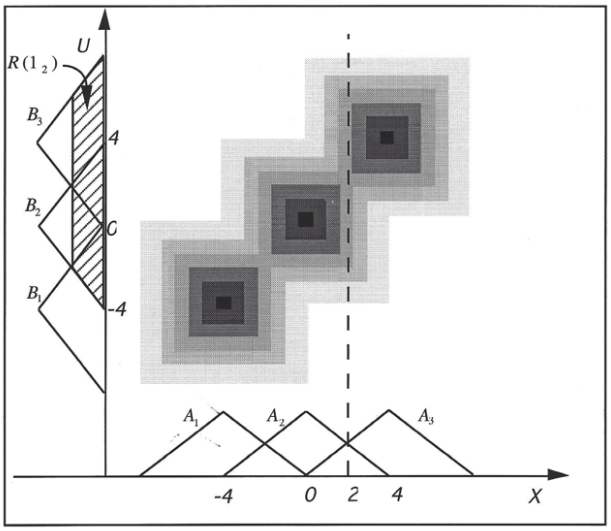

Figure 3: (View in browser) Fuzzy graph expressed by 3 rules: if A_1 then B_1, if A_2 then B_2, if A_3 then B_3. The fuzzy graph is the union of the Cartesian product of corresponding antecedents and conclusions, $i.e. A_1 \times B_1, A_2 \times B_2$ and $A_3 \times B_3$.The fuzzy output is obtained by taking the Cartesian product of the graph with I_2, where I_2 represents the fuzzy set with membership grade 1 at $x = 2$ and 0 everywhere else. This results in the output $R(I_2)$ for $x = 2$.

Related to this viewpoint is the alternative inference mechanism based on fuzzy valued interpolation3 (see Figure 4). In order to respect the meaning of the fuzzy conclusion, being a fuzzy restriction on the possible control actions, it is important to consider the fuzzy mapping obtained before the defuzzification step is applied (see Figure 1). As many rule bases for fuzzy controllers have been synthesized based on learning algorithms such as genetic algorithms and neural network backpropagation like-algorithms which focused on getting the mapping right from crisp input to crisp output, the semantics of the fuzzy conclusions was

completely overlooked, reducing it to a second class object.

However, this does not mean we should abandon fuzzy controllers. By preforming learning before the defuzzification phase, the semantics of a fuzzy conclusion can be respected, allowing to obtain a richer model of the problem where the width of a fuzzy conclusion is no longer second class object and where the use of fuzzy rules truly makes sense. It for example allows to learn a model of the system where the fuzzy conclusions express a fuzzy restriction on the control actions, resulting a rule-base that can be used as a fuzzy controller. Doing so, fuzzy controllers are an interesting formalism in their own right.

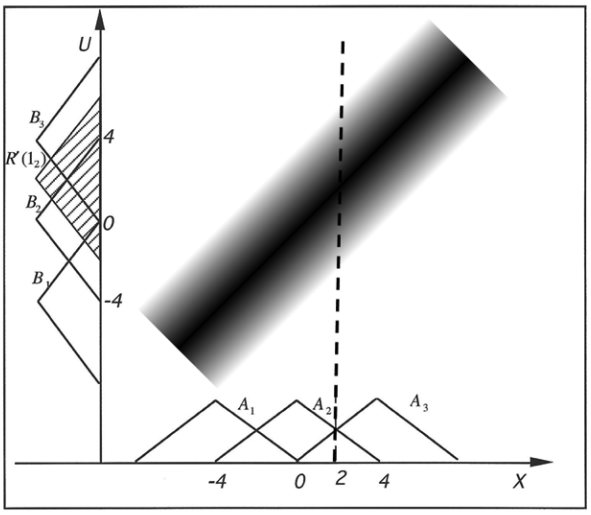

Figure 4: (View in browser) Fuzzy Lagrange interpolation. Here the fuzzy antecedents A_1, A_2 and A_3 are interpreted as a piecewise linear interpolation net. The fuzzy numbers corresponding to the inputs -4, 0 and 4 are respectively B_1, B_2 and B_3. To obtain the interpolating fuzzy-valued function at $x = 2$, the 2 fuzzy numbers B_2 and B_3 are combined using a weighted sum, resulting in $R'(I_2)$. This figure provides only an approximate representation.

References

L.A. Zadeh, Fuzzy sets, *Information and Control*, Volume 8, Issue 3, 1965, Pages 338-353, ISSN 0019-9958, http://dx.doi.org/10.1016/S0019-9958(65)90241-X.

L.A. Zadeh, The concept of a linguistic variable and its application to approximate reasoning-III, *Information Sciences*, Volume 9, Issue 1, 1975, Pages 43-80, ISSN 0020-0255, http://dx.doi.org/10.1016/0020-0255(75)90017-1.

Interview with Pouyan Ziafati of LuxAI

by Marc van Zee

Could you tell us something about your background?
My name is Pouyan Ziafati, my background is in robotics and AI. I did a master in advanced robotics, funded by the European Union, coordinated by the Warsaw University of Technology and the University of Genoa. I did my PhD in the area of robotic software and knowledge representation and reasoning, focusing on decision making for robotics at the University of Luxembourg and Utrecht University. After this, I became the principle investigator of a Proof of Concept (POC) project at the University of Luxembourg, where the aim was to use the research results of my PhD and turn it into an application. The application we chose was social robotics.

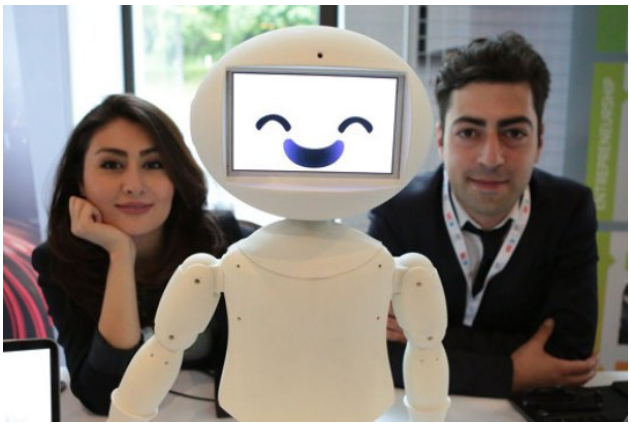

Aida Nazarikhorram and Pouyan Ziafati of LuxAI

What is the vision of LuxAI?
Our vision is to make social robots usable for everybody. For social robotics, there are currently applications in health and education. In these applications, you need a customized robot applications for individuals. It is very important to us that teachers and therapists can use robots. We see very promising results in these applications, such as therapy for autistic children, or rehabilitation of elderly people. However, this is mostly research, and we want actually enable real people to use these robots. We want to make social robotics economic and user friendly.

In an interview with Manuela Veloso (this newsletter), she said that one condition for robots to enter our lives is that we should allow robots to make mistakes. Do you think this as well?
I think this very much depends on the application domain. In the domain where our robots are developed (health and education), there are many questions about who is going to pay if something goes wrong. So as a company developing robots, we should try to minimize the risk of robots making mistakes, and we should be very clear about our responsibilities. We have to minimize the risk of anything going wrong. I do understand the logic behind Professor Veloso's argument: At some point there is no definition of perfection. When you think about communication between human and robots, it's very difficult to say what is really wrong and what is right.

Do you work on verbal communication as well?
Currently not. In most of our applications we are looking for others ways of communication that are more reliable and robust. For instance, we use vision to recognize markers and pictures.

What would your robot do in the teaching domain?
The main application that we build our robot for is to assist autistic children. They often have problems with communication and social interaction. They also have problems understanding facial expression. Our

robot helps those children to understand facial expressions and emotions. For instance, simply by explaining what is a happy or sad face. Our robot can simplify these expressions, but also exaggerate them. Our robot can tell social stories as well. It can teach children what are good behavior and what to say in each context. For instance, when to say hello, to share toys, and when to say thank you. For each of these stories it then shows emotions that belong to them. You also see in the picture (see picture above, *red.*) that the robot has a big screen to show faces. This is on purpose, in order for the robot to be very expressive for children with developmental problems.

How many robots do you have, how big is your company, and are you planning to expand?
We build a prototype of our robot and we did initial tests with teachers and therapists in autism centers. We are confident that they can easily work with our application and they can make customized robot applications. They also were positive towards using the robots. We are currently manufacturing the robots and we hope to sell the first version to researchers by the end of this year, and then produce a larger number next year. We are also planning a large scale experiment to show the effectiveness of our robots for autism therapy. Finally, we are working on creating applications for teaching foreign languages to children in kinder garden.

You have a company with your wife Aida together. This is quite unique, how did this happen?
Aida is a medical doctor. In Iran, she created an NGO for children with diabetics. When she came to Luxembourg she started working at a start up and she liked the idea of having a start up herself. We started thinking together how to do something at the intersection of robotics and health. We first thought of making applications or AI tools for diagnosis, but we finally settled on the idea of building social robots. Aida has a lot of knowledge in medical application, and she can communicate with experts in those domains (e.g., expert in autism or Geriatric Medicine). Therefore, she is managing the collaboration with our partners. She is also a fantastic presenter.

Would you be interested in collaborating with universities in the Benelux as well?
Yes, this is currently on of our focuses. We would like to initiate as many collaborations as possible for different types or social robot application. We are mostly interested in autism and Geriatric Medicine, in particular, post-stroke rehabilitation. If someone is interested to try our robot we are very open to discuss about this.

It sounds like the BNAIC conference would be interesting for you!
Yes indeed, we will definitely try to be there and bring on of our robots!

Do you think Luxembourg is a good place for a start up?
Yes, I think it is a heaven for start ups. The country is investing a lot in it so there are many means for support. There are several incubators providing different types of support, for instance cheaper offices or mentors that help you developing business plans. LuxInnovation [6] is a fantastic place with a big network that you can use to find possible collaborators. FNR, the national research fund of Luxembourg[7], is great for funding and has the POC project, which is perfect to develop a application from your research.

Do you think having a PhD is helping you in setting up a company? Not just in terms of expertise, but also in

[6] See `http://www.luxinnovation.lu/`
[7] See `http://fnr.lu/`

terms of project management.
Yes I think so. I learned how to organize and respect deadlines. I also learned how to read fast and write professionally, and how to communicate.

Would you ever consider going back to research?
Currently business is very exciting, but I also loved the time I did research. If I had extra time I would still like to do some research. The nice thing about my current position is that I am still involved in research, but I am managing it more than that I am doing it myself. I always had an eye for application in my research, and I think having a direction helps to narrow the space to investigate.

Where do you see LuxAI in five years?
In five years we are one of the biggest social robotics players in the world.

AI: From Science Fiction to Science Non-Fiction (Column)

by Nieske Vergunst

When I started studying Cognitive Artificial Intelligence in 1999, the term was mostly met with laughter and confusion. “Artificial intelligence? Do you really need that? You’re such a smart girl already!” Especially relatives and acquaintances of the baby boomer generation weren’t quite up to speed. The term ’AI’ was known mostly as an acronym for ’artificial insemination’. I kept having to explain that making computers smart was a really interesting challenge. Mention of Deep Blue, who had won its legendary game against Garry Kasparov two years prior, didn’t help much. Chess was for nerds, after all, not to mention things with chess and computers.

A lot has changed in the past 17 years. According to media database LexisNexis, 2015 saw the publication of almost 100 articles containing the keyword ’robot’ in Dutch major newspapers. The boom seems to have just begun: 237 news articles about robots were published in the year 2016 so far – and it’s only September. In comparison: in 2011, almost 50 news articles about robots were published, and in the years 2012 to 2014, approximately 80-90 per year.

While there are plenty of reports about killer robots and war drones, Dutch media attention about robots in 2016 seems to be mostly related to robots moving from academia into society: robots in the classroom, robots in museums, robots that deliver pizza, pour beers, and do the dishes[8]. It seems to be especially the service robot that is interesting for a broad audience. See also the Roomba, which vacuums floors in millions of households[9] and has brought its mother company a half billion dollar profit in 2015[10].

The friendly, handy household robot seems to be in sharp contrast to the average movie robot, who is more of a fighter than a helper. After the 1919 movie The Master Mystery[11], the first movie starring a robot avant la lettre (then called ’The Automaton’), more and more movies were made about robots, with an absolute high in the 1980’s, the decade of Blade Runner, The Terminator, RoboCop, Transformers, and plenty of Japanese anime titles starring robots. In the 1980’s, 0.16% of all movies featured robots.[12] That may not seem much, but it’s only about ten times less than movies tagged ’murder’. Whereas the absolute number of movies about robots in subsequent decades is larger, they do make up a smaller percentage of the total amount of movies released (less than 0.03% of 2015 movies, for example).

Though robot movies are tagged mostly with keywords like ’death’, ’fight’, and ’murder’, recent movies seem[13] to be moving away from robot violence. Whereas 20th century movies seem to star mostly fighting robots – of course excepting friendly R2-D2 and C-3PO, but the counterexamples are few and far between – current movie robots appear to be less about fighting and more about thinking. Recent years have seen some very successful movies starring robots that reach their goals by other means than violence, such as *WALL-E, Robot & Frank, Interstellar, Ex Machina, Big Hero 6,* and the aptly titled films *Robots* and *Artificial*

[8] http://academic.lexisnexis.nl/
[9] http://www.irobot.com/About-iRobot/Company-Information/History.aspx
[10] http://www.irobot.com/About-iRobot/Company-Information.aspx
[11] https://en.wikipedia.org/wiki/The_Master_Mystery
[12] http://www.imdb.com/search/keyword
[13] For lack of time, I’m afraid I have to spare you the statistical details.

Intelligence: AI, to name a few.

In the meantime, the fear of robots seems to have moved from the abstract ("they're going to kill me!") to the concrete: will a robot take my job? After immigrant workers and the economic crisis, the robot seems to be the next threat to the workforce. Even though robots are commonly overestimated – I have yet to see the first convincing argument that robots are going to take over the world and kill all humans – there are, of course, justified fears.

Robots will surely have an impact on the workforce, just like the invention of the assembly line, electricity, the washing machine, the computer, and so on. For you and me, chances are that we'll be able to keep our jobs. According to a study by researchers at Oxford University and Deloitte[14], physical scientists have a 20% likelihood of being replaced by robots. For me, depending on whether I call myself an author or a journalist – as a science writer, both are partially true – that chance is 33% or only 8%, respectively.

The coming years and decades will probably see robots and AI moving more into the ubiquitous, but like with all developments, it's difficult to predict the details. Similarly to the internet a few decades earlier, we are the lucky people who are able to witness that development. Whatever happens, and whether or not our work will be outsourced to smart robots, at least when we say that we've worked in AI and robotics, people won't be confused anymore. Whether they'll be impressed or angry... That's something only time will tell.

[14] `http://www.bbc.com/news/technology-34066941`

The Robot Baby Project

by A.E. Eiben

In May 2016 the team of VU scientists in the figure on the right presented the first 'robot baby'. Popular media covered the presentation extensively, but the corresponding scientific paper is still under review. In this note I briefly summarize what the Robot Baby Project (RBP) did and what this means.

This project is part of a greater endeavor towards the Evolution of Things, as outlined recently in Nature (Eiben and Smith, 2015). In particular, the RBP is a proof of concept project to demonstrate that robots can reproduce. To achieve this goal, the project adopted two premises. The first one is the use of the Triangle of Life system model (Eiben et al., 2013); the second one is simplification.

The Triangle of Life is a generic system architecture that decomposes a physical evolutionary system and forms a blueprint for implementation. In essence, it specifies the basic life cycle that does not run from birth to death, but from conception (being conceived) to conception (conceiving one or more children) and it is repeated over and over again, thus creating consecutive generations of robots. A real world implementation will result in an evolving population of robotic organisms, where the bodies as well as the brains can adapt to the given environment.

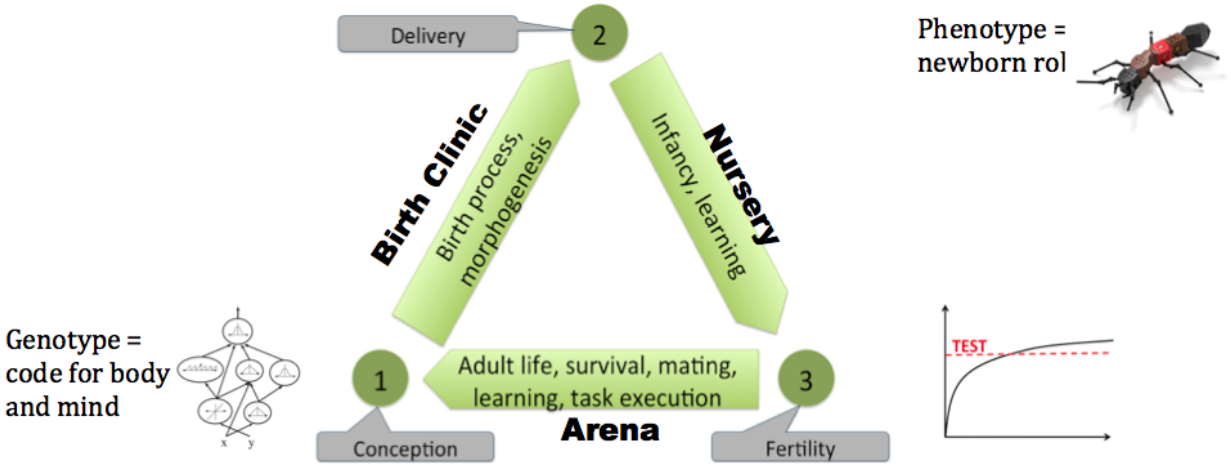

Diagram of a fully physical artificial evolution system based on the Triangle of Life architecture. The diagram captures a workflow not from birth to death, but from being conceived to conceiving a child. The pivotal moments that span the triangle are: Conception (a new genotype is created, construction of a new robot starts), Delivery (construction of the new robot is completed), and Fertility (a robot becomes an adult).

As the name suggests there three principal stages in a life cycle: Morphogenesis, Infancy, and Adult Life and three corresponding system components: the Birth Clinic, the Nursery, and the Arena. The Birth Clinic constructs robot phenotypes based on a given genotype. Newborn robots start in the Nursery where they are trained and tested. If they pass the test after training they are declared fertile and enter the main arena; otherwise they are removed and recycled. The Nursery increases the chances of success in the Arena and prevents reproduction of inferior robots. The Arena is the place where the robots have to live and work: survive, reproduce, learn, and perform user-defined tasks. Mate selection can be autonomous and/or human controlled (breeding). Selected parents transmit their genome to the Birth Clinic that performs crossover and mutation and constructs the child. It is important to note that evolution results from the whole life cycle, while leaning takes place in two of the stages/components: during Infancy (Nursery) and Adult Life (Arena).

Simplification being our second main premise, we limited the set of body parts to a minimum. We used three types of 3D-printable components: a head block, a body block, and a joint. To obtain a functional robot we combined these with not printable parts, such as servos to drive the joints, light sensors to be the 'eyes', a Raspberry Pi to be the 'brain', and a rechargeable battery to provide energy. These components and their possible combinations determine all realizable robot bodies that can be created within our system.

We decided to make the first parent with the shape of a 'spider', be it with fewer legs. It had a central module with the CPU, the battery, and a light sensor in the middle and four limbs with two blocks and two servo driven joints. The color of the 3D-printed blocks was blue. The second parent had the shape of a 'gecko'. Its head, the central module with the CPU, the battery, and a light sensor, was in front. Its body showed a left-right symmetry, the limbs are shorter than the spider's. The color of the 3D-printed blocks was green.

The two parents went through the infancy stage and became adults. To this end, they had to learn locomotion and navigation to a specific spot, the 'mating corner' of the habitat, supervised by an overseeing camera that provided feedback on the robots' behavior. Once they met in the mating corner, they mated (virtually) and sent their genome by WiFi to the Birth Clinic. This consisted of a computer, a 3D printer, and the collection of organs. The parental genotypes were randomly recombined into a new one and a new robot was printed and hand-assembled according to this specification. This delivered the first 'robot baby' and concluded the robotic life cycle.

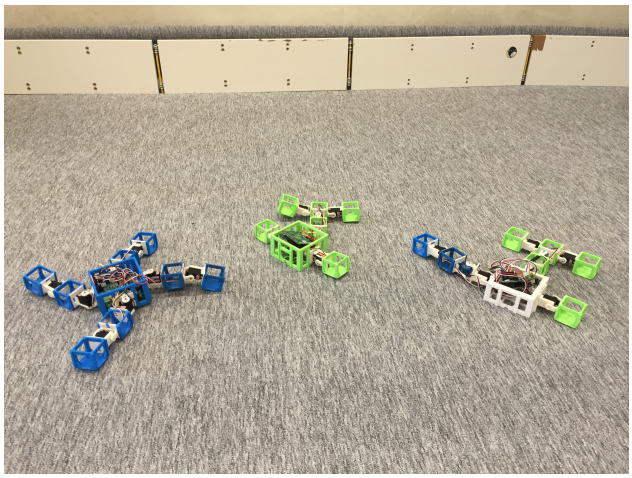

The RBP achieved its objective and demonstrated that robots can reproduce. Strictly speaking we did not demonstrate that robots can evolve, because one child does not make a generation, and even if it did (extreme case with n=1) one generation does not make evolution. However, "by induction" we did prove the possibility of creating many consecutive generations. To get evolution we then only need selection, which is arguably the easy part of the equation. Environmental selection is for free in a real world setting (unfit robots break or 'starve') and mate selection is straightforward: the corresponding procedure can be based on any user defined preference and/or measurable fitness indicators, such as the level of energy, the number of previously conceived children, or the age of the robot.

The RPB designates a milestone as it represents a new form of evolutionary robotics (fully in hardware) and a holistic design approach where morphologies and controllers of robots are not separated and can be developed together in the real environment. On the long term, it seeds a disruptive robotic technology where robots are not designed and manufactured in the traditional way, but 'develop themselves' on the job through evolution and learning. For scientific fields concerned with artificial evolution (evolutionary computing and artificial life), it initiates a paradigm change in by changing the substrate from digital to physical. This will provide new insights into the working of evolutionary systems. Last, but not least, evolving robot systems represent a new breed of machines that can change their form and behavior not in error but by design.

References

A.E. Eiben and J. Smith, From evolutionary computation to the evolution of things, *Nature*, 521:476-482, 2015.

A.E. Eiben, N. Bredeche, M. Hoogendoorn, J. Stradner, J. Timmis, A.M. Tyrrell and A. Winfield, The Triangle of Life: Evolving Robots in Real-time and Real-space, P. Lio et al. (Eds.), *Proceedings of the 12th European Conference on the Synthesis and Simulation of Living Systems (ECAL 2013)*, MIT Press, 2013, pp. 1056-1063.

Event Reports

The Future of Robot Rescue Simulation workshop

by Arnoud Visser

This spring a workshop was held in the Lorentz Center in Leiden on the Future of Robot Rescue Simulation. This was not a regular Lorentz workshop, because the preparations of a regular Lorentz workshop start 1½ year in advance. In our case, a new roadmap and rulebook had to be prepared before the next RoboCup competition, which meant that the workshop had to be organized in short notice (half a year). Still, the workshop was completely in line with the philosophy of the Lorentz Center: to be a home for creative researchers and initiate new collaborations and interactions between scientists from different countries with varying seniority. At the end, the selected participants came from the Netherlands, Turkey, Italy, Tunisia, Austria, United Kingdom, Switzerland, Iran, Germany, Portugal, Peru, Malaysia, and Japan. They had participated in competitions as the DARPA Challenge, Japanese Virtual Robot Challenge, euRathlon, UAE Drones for Good, RoboCup Junior, RoboRace, Mid-size Sumo, and the RoboPoly challenges.

Figure 5: Participants of the Future of Robot Rescue Simulation workshop

The international community was very happy with the warm welcome they received in The Netherlands, which was partly due to the Wine & Cheese party sponsored by the KNVKI. Other social events were the city tour through Leiden and the conference diner in the Faculty Club, sponsored by respectively The Construct and Mathworks.

See `https://staff.fnwi.uva.nl/a.visser/activities/FutureOfRescue/`

ACAIS 2016: Best of ACAIS

by Symposium Commissie CognAC

On November 5-6, 2015, the 27th Benelux Conference on Artificial Intelligence (BNAIC'15) was held in Hasselt, Belgium, in a renovated old prison currently used by the Faculty of Law. Members from the Business Informatics research group at Uhasselt were responsible this conference. As usual, BNAIC'15 was organized under the auspices of the Benelux Association for Artificial Intelligence (AIABN-BNVKI) and the Dutch Research School for Information and Knowledge Systems (SIKS). The conference was a lively mixture of 15 oral presentation sessions in three parallel tracks, a poster and demo session, plus two invited speakers. On Thursday Dr. Elpiniki Papageorgiou started the conference with a keynote titled "Fuzzy Cognitive Maps: Methods and Applications". Dr. Elpiniki Papageorgiou is an Assistant Professor in the Department of Computer Engineering at Technological Educational Institute, University of Applied Sciences of Central Greece, Lamia, Hellas. She has been working in the area of computational intelligence for over fourteen years conducting research mainly in fuzzy cognitive maps. More explicitly, she has been working as researcher in several research projects related with the development of novel computational intelligence methodologies for decision support systems, intelligent algorithms for decision making, data analysis and mining, knowledge-based systems and expert systems. In the keynote she highlighted the core methodology of fuzzy cognitive map modeling and inference, their dynamic characteristics and learning capabilities, some promising extensions and their evolutionary structures capable to solve complex modeling and decision making problems. Selected innovative applications of them to diverse research domains during the last years were presented to show their usefulness in modeling complex problems. Next the first six oral presentation sessions , the demo session and the poster session were responsible for a heavily loaded day. The demo sessions were really impressive . Coffee breaks and a small but delightful lunch were offered and appreciated . The first day was concluded by a guided city walk in Hasselt , the city of taste , and by a conference dinner in "'t Klein Genoegen" and for some by a strong Belgian beer. Participants confirmed the 'City of Taste' label. On the second day of the conference, the second invited speaker opened the program. Dr. Stephan Onggo is a lecturer (Assistant Professor) in the Department of Management Science at Lancaster University Management School (LUMS), Lancaster, United Kingdom. His research interests are in the areas of simulation methodology (Conceptual Modelling, Modelling Paradigms including discrete-event, system dynamics, and agent-based), simulation technology (Parallel/Distributed Simulation and cloud-based Simulation), simulation applications (healthcare, public sector, supply-chain), and business process modelling (BPMN). In his keynote he introduced Agent-based simulation (ABS) and its widely use in social sciences and operational research. ABS was born from the field of Artificial Intelligence (AI), especially in the subfields of multi-agent systems, artificial life and distributed AI. Dr. Onggo presented the use of ABS to represent intelligent behaviors in a society, he identified AI components in ABS research and highlighted potential collaborative works between ABS communities and AI communities.

After the coffee break, the conference participants attended nine oral presentation sessions . The presentations were in general of high quality. The annual BNVKI meeting was also held. The conference was concluded by short award closing session. Koen Vanhoof presented the award for the best paper to Vincent Nys, Jon Sneyers en Daniel De Schreye for their paper Automatic music teaching : a logic-probabilistic étude generator. After that, Jaap van den Herik awarded the SKBS best demo award to Wiebe Van Ranst and Joost Vennekens for their demo Ultra-low-latency endoscopic image stabilization. Finally, participants thanked the organizers and the location of BNAIC 2017, which will be held in Amsterdam co-organized by both universities of Amsterdam, was announced. We wish to thank all the researchers, the participants, the keynote speakers, the Program Committee and additional reviewers for their careful paper reviews;

the BNVKI board for their advice; and the student volunteers for their invaluable help in making this conference come true

Highlights of EurAI General Assembly

Three EurAI board members completed their terms at ECAI:

- Mike Wooldridge
- Jerome Lang
- Ann Nowe

To replace them, three new board members were elected:

- Hector Geffner (Spain)
- Catholijn Jonker (Netherlands)
- Steven Schockaert (Belgium/UK - TREASURER)

In addition, a new chair and deputy chair of the EurAI boardwere elected:

Gerhard Lakemeyer, (Germany), elected chair 2016-18
gerhard@informatik.rwth-aachen.de

Barry O'Sullivan (Ireland), elected deputy 2016-18 b.osullivan@cs.ucc.ie

Event Announcements

BNAIC 2016 (Amsterdam)

The 28th Benelux conference on Artificial Intelligence (BNAIC 2016) will be jointly organized by the University of Amsterdam and the Vrije Universiteit Amsterdam, under the auspices of BNVKI and SIKS. BNAIC 2016 will be held in Hotel Casa 400, Amsterdam, The Netherlands, on Thursday 10 and Friday 11 November, 2016.

BNAIC 2016 is promising to be an interesting and varied congress, in which all different types of Artificial Intelligence are represented.

Invited Speakers

- Hado van Hasselt (Google DeepMind)
- Manuela Veloso (Carnegie Mellon University)
- Marc Cavazza (University of Kent)

TEdX special session There is a special session this year with presentations, in which among others Catholijn Jonker (TU Delft) and Leon van der Torre (University of Luxembourg) will speak.

'Research meets Business' event

Panel discussion about 'Societal Implications of Social Robots', with cooperation of Elly Konijn (VU), Mark Neerincx (TNO/TU Delft) and Aimee van Wynsberghe (UTwente).

More information: http://www.bnaic2016.org/.

BNVKI Membership Fees

In the table below you can find the BNVKI membership fees.

	2015
Regular members	€ 20,-
PhD students	€ 10,-
Master students	€ 10,-

Table 1: BNVKI Registration Fees

Becoming a BNVKI member makes you automatically an ECCAI member and allows you register at a reduced registration rate for certain major events, such as ECAI and ACAI. By increasing the number of BNVKI members, our AI community can also nominate more colleagues to become ECCAI fellows, as the maximum number of fellows we are allowed to have is proportional to the number of members. Finally, it might be good to know that ECCAI has decided to sponsor international events through invited speakers and these invited speakers need to be an ECCAI member over the past years.

If you want to know where our members are currently located, check out http://wilma.vub.ac.be/ dvandeun/mapje.html, if your affiliation is not represented, or you would like to see a larger dot, become a member and convince you colleagues to join as well.

Board Members BNVKI

How to Subscribe?

The BNVKI-AIABN Newsletter is a direct benefit of membership of the BAIAI: Benelux Association for Artificial Intelligence. Membership dues are € 20 for regular members and € 10 for students (AIO’s or master). In addition, members will receive access to the electronic version of the European journal AI Communications. The Newsletter appears quarterly. For more information, please visit our website and go to “Membership and Benefits”.

Copy

The editorial board welcomes product announcements, book reviews, product reviews, overviews of AI education, AI research in business, and interviews. Contributions stating controversial opinions or otherwise stimulating discussions are highly encouraged. Please send your submission by E-mail (MS Word or text) to board@bnvki.org.

Advertising

It is possible to have your advertisement included in the BNVKI/AIABN Newsletter. For further information about pricing etc., see our website, section “Sponsoring Rules”.