AI Profiles: An Interview with Kristian Kersting

Marion Neumann (Washington University in St. Louis; m.neumann@wustl.edu)
DOI: 10.1145/3284751.3284758

Introduction

This column is the sixth in our series profiling senior AI researchers. This month we interview Kristian Kersting, Professor in Computer Science and Deputy Director of the Centre for Cognitive Science at the Technical University of Darmstadt, Germany.

Figure 1: Kristian Kersting

Biography

After receiving his Ph.D. from the University of Freiburg in 2006, he was with the MIT, Fraunhofer IAIS, the University of Bonn, and the TU Dortmund University. His main research interests are statistical relational artificial intelligence (AI), probabilistic deep programming, and machine learning. Kristian has published over 170 peer-reviewed technical papers and co-authored a book on statistical relational AI. He received the European Association for Artificial Intelligence (EurAI, formerly ECCAI) Dissertation Award 2006 for the best AI dissertation in Europe and two best-paper awards (ECML 2006, AIIDE 2015). He gave several tutorials at top AI conferences, co-chaired several international workshops, and cofounded the international workshop series on Statistical Relational AI (StarAI). He regularly serves on the PC (often at senior level) for several top conference and co-chaired the PC of ECML PKDD 2013 and UAI 2017. He is the Speciality Editor in Chief for Machine Learning and AI of Frontiers in Big Data, and is/was an action editor of TPAMI, JAIR, AIJ, DAMI, and MLJ.

Getting to Know Kristian Kersting

When and how did you become interested in AI?

As a student, I was attending an AI course of Bernhard Nebel at the University of Freiburg. This was the first time I dived deep into AI. However, my interest in AI was probably triggered earlier. Around the age of 16, I think, I was reading about AI in some popular science magazines. I did not get all the details, but I was fascinated.

What professional achievement are you most proud of?

We were collaborating with biologists on understanding better how plants react to (a)biotic stress using machine learning to analyze hyperspectral images. We got quite encouraging results. The first submission to a journal, however, got rejected. As you can imagine, I was disappointed. One of the biologists from our team looked at me and said ”Kristian, do not worry, your research helped us a lot.” This made me proud. But also the joint work with Martin Mladenov on compressing linear and quadratic programs using fractional automorpishms. This provides optimization flags for ML and AI compilers. Turning them on makes the compilers attempt to reduce the solver costs, making ML and AI automatically faster.

What would you have chosen as your career if you hadn’t gone into CS?

Physics, I guess, but back then I did not see any other option than Computer Science.

What do you wish you had known as a Ph.D. student or early researcher?

That "sleep is for post-docs," as Michael Littman once said.

Artificial Intelligence = Machine Learning. What's wrong with this equation?

Machine Learning (ML) and Artificial Intelligence (AI) are indeed similar, but not quite the same. AI is about problem solving, reasoning, and learning in general. To keep it simple, if you can write a very clever program that shows intelligent behavior, it can be AI. But unless the program is automatically learned from data, it is not ML. The easiest way to think of their relationship is to visualize them as concentric circles with AI first and ML sitting inside (with deep learning fitting inside both), since ML also requires to write programs, namely, of the learning process. The crucial point is that they share the idea of using computation as the language for intelligent behavior.

As you experienced AI research and education in the US and in Europe, what are the biggest differences between the two systems and what can we learn from each other?

If you present a new idea, US people will usually respond with "Sounds great, let's do it!", while the typical German reply is "This won't work because ...". Here, AI is no exception. It is much more critically received in Germany than in the US. However, this also provides research opportunities such as transparent, fair and explainable AI. Generally, over the past three decades, academia and industry have been converging philosophically and physically much more in the US than in Germany. This facilitate the transfer of AI knowledge via well-trained, constantly learning AI experts, who can then continuously create new ideas within the company/university and keep pace with the AI development. To foster AI research and education, the department structure and tenure-track system common in the US is beneficial. On the other hand, Germany is offering access to free higher education to all students, regardless of their origin. AI has no borders. We have to take it from the ivory towers and make it accessible for all.

What is the most interesting project you are currently involved with?

Deep learning has made striking advances in enabling computers to perform tasks like recognizing faces or objects, but it does not show the general, flexible intelligence that lets people solve problems without being specially trained to do so. Thus, it is time to boost its IQ. Currently, we are working on deep learning approaches based on sum-product networks and other arithmetic circuits that explicitly quantify uncertainty. Together with colleagues—also from the Centre of Cognitive Science—we combining the resulting probabilistic deep learning with probabilistic (logical) programming languages. If successful, this would be a big step forward in programming languages, machine learning and AI.

AI is grown up - it's time to make use of it for good. Which real-world problem would you like to see solved by AI in the future?

Due to climate change, population growth and food security concerns the world has to seek more innovative approaches to protecting and improving crop yield. AI should play a major role here. Next to feeding a hungry world, AI should aim to help eradicate disease and poverty.

We currently observe many promising and exciting advances in using AI in education, going beyond automating Piazza answering, how should we *make use of AI to teach AI*?

AI can be seen as an expanding and evolving network of ideas, scholars, papers, codes and showcases. Can machines read this data? We should establish the "AI Genome", a dataset, a knowledge base, an ongoing effort to learn and reason about AI problems, concepts, algorithms, and experiments. This would not only help to curate and personalize the learning experience but also to meet the challenges of reproducible AI research. It would make AI truly accessible for all.

What is your favorite AI-related movie or book and why?

"Ex Machina" because the Turing test is shaping its plot. It makes me think about current

real-life systems that give the impression that they pass the test. However, I think AI is hard than many people think.

Help us determine who should be in the AI Matters spotlight!

If you have suggestions for who we should profile next, please feel free to contact us via email at aimatters@sigai.acm.org.