CV

Frank van Harmelen

29-10-1960, Bussum

- **1978-1983** mathematics and computer science studies, University of Amsterdam
- **1985-1989** PhD research into meta-level reasoning, University of Edinburgh
- **1990-1995** postgraduate AI researcher, University of Amsterdam
- **1995-2001** AI researcher with focus on semantic web and knowledge representation, VU University Amsterdam
- **2001-present** professor of Knowledge Representation and Reasoning, VU University Amsterdam
- **2019-present** principal scientist Hybrid Intelligence: augmenting human intellect, alliance of six Dutch universities, winner Gravitation subsidy

'We want AI to be a fully-fledged colleague'

Frank van Harmelen is one of the initiators of the largest AI research project in Dutch science history: Hybrid Intelligence. His mission: to develop computers that can make plans and solve problems.

Text: Tamara van 't Woud
Photography: Bob Bronshoff

You are Professor of Knowledge Representation and Reasoning. What makes reasoning difficult for a computer?
'Reasoning requires knowledge. How do we represent knowledge in such a way that a computer can reason with it? Things that are very easy for people are very difficult for a computer – and vice versa. In AI, we call this Moravec's paradox, named after a researcher who observed this back in the 1980s.

A standard example of this is having a conversation in a natural way: hearing what you say, understanding the language, understanding the context so well that you give it the correct meaning, and then giving an answer that is tailored to the listener. I talk to you differently than I talk to my colleagues. I know pretty much what you know, and I tune my answers to that. The conversation has a purpose. And I know that you know that I know what the purpose is. A toddler of four can assess that as well. But for a computer, that's very difficult.'

That sounds very practical, that people are good at what computers are bad at and vice versa. Why are we afraid that computers will replace people?
'The basic idea underlying the Hybrid Intelligence project is precisely that people and computers can complement each other. We want to work towards a hybrid team of people and computers working together. We think such a team is better at some tasks than a team of just people or just computers.

Within the HI project, everyone contributes on a different level. In Delft, for example, they are good at the social side of computers, so they work on the collaborative aspect. A colleague from the VU is working on langua-

ge technology, which is about the communication between computers and humans. In Groningen, there is a team member who works together with cognitive scientists. And a UvA colleague is working on self-learning computers. Each participating university has its own role. The two universities in Amsterdam are leading the project, and the other four are contributing expertise that we don't have.'

What do you expect from the future of AI?
'I think we're dropping the science fiction film-like idea of computers being intelligent in the same way that people are, and we are focusing more and more on human-computer collaboration. To do that, we have to study other AI questions than we've done up to now.'

Such as?
'Computers are a little autistic now. People are constantly thinking about each other's knowledge, objectives, and motivations. This enables us to work well together in teams. Computers have to get better at this. So far, this is an underexposed aspect of AI, which is why it is so good that six Dutch universities are going to work on this in a collaborative venture.'

You sketch an optimistic scenario, but the world consists of more than well-meaning scientists. What is the pessimistic scenario?
'Uneven distribution of the benefits of technological revolutions in society. We saw the same thing happen with the Industrial Revolution: initially factory owners became very rich, and the workers became very poor. It took a hundred years before the benefits were distributed a little more evenly. There are concerns that this will be the case with AI.'

How do the results of the Gravitation Project reach society?
'Our particular focus is on application in education and healthcare. For example, a colleague is experimenting with robots in children's hospitals. They can provide information if a child asks when mom and dad will be visiting, or why they need to be in hospital. A robot like that would be part of the hospital staff.

The challenge for us is to become the first team to publish a scientific article of which AI – which we have already christened AI-nstein – is a co-author. That means AI must have contributed to every step of the process, from the formulation of the hypothesis to the final publication. We want AI to join the team as a fully-fledged colleague.'

Including lunch breaks?
'That could be the hardest thing of all.' ■

column

Fundamental centre of development

VU University and its (Amsterdam) partners are making dizzying progress in various fields of science using artificial intelligence (AI) and data science. What I find striking about our research, to which more than twenty VU professors and their research groups are contributing, are three things.

First of all, on the basis of a fundamental structure, our research is making astonishing contributions to many applications. In the field of health and care, for example, we are developing care robots, AI for the prevention of suicide, and the development of AI-driven therapies to support people with depression. With regard to (business) organisations, our research ranges from the development of an instrument for the police that predicts misconduct and crime to an AI instrument to manage successful recruitment and selection at companies.

A second notable feature of our research is that it often focuses on hybrid intelligence, that is how artificial and human intelligence complement each other. The aim is optimal collaboration between people and machines instead of replacing people by machines. Intended applications here are the use of robots as teaching assistants in the classroom, robots as buddies for children in hospitals, and AI as a collaborative partner of scientists.

Finally, it strikes me that we are working together so much in this field of science. Not only with other knowledge partners, in Amsterdam and internationally, but also with other organisations on a large scale. And research teams are not only made up of hard scientists, but also scientists from other disciplines, such as linguists, business experts, doctors, psychologists, criminologists, and human movement scientists.

All in all, I am enormously proud of the rapid growth of our research and education in this field. Although this sometimes makes a new building too small at its opening, our contribution to social challenges based on often multidisciplinary and translational research is spectacular. It is, therefore, logical that we see artificial intelligence and data as a fundamental centre of development of the VU University Amsterdam.

Mirjam van Praag
President of the Vrije Universiteit Amsterdam

Intensive **intelligence**

A collaboration of science, medical expertise, and business: you don't see that much within AI. A group within the AI coalition is trying to incorporate data from ICU patients into models in such a way that doctors can make better decisions about treatment.

Text: Marleen Hoebe

The data of patients who have received intensive care (ICU) treatment is very valuable for new ICU patients.

By looking at large amounts of data, it is easier to determine in advance which treatment will work best for a patient. A group within the AI coalition is trying to incorporate ICU patient data into models in such a way that doctors can make better decisions about treatment.

One of the collaborations within AI technology for people is the collaboration between intensivist Paul Elbers, AI scientist Mark Hoogendoorn, and Pacmed founder Willem Herter. All three want to apply artificial intelligence (AI) in healthcare, but they all look at it from a slightly different perspective. Together, they connect science, medical expertise, and business. That is something you hardly see within the area of AI. Why is that the case? And what is the reason that these men decided to bring together a faculty, a medical hospital, and a business?

'It's very complicated to make a model truly implementable'

Paul Elbers

Intensivist, Amsterdam UMC, VUmc location, Leader of the collaboration

'The data of ICU patients forms the basis for our project. Using this data, we can use AI to develop models that predict, for example, how likely it is that a patient discharged from the ICU will eventually have to return to the ICU. We already have a readmission model and it is available in the hospital – but only for research purposes. We hope to be able to implement it quickly in the clinic so that the model can support doctors in the decision to dismiss someone from the ICU sooner if it is safe to do so or later if necessary. This can increase both the quality and the capacity of the ICU.

Other models that we are now creating examine which treatments give the best results. For example, which medication works best, or which manner of artificial respiration is better. And we have many other models that are in the early stages of development. We're not the only ones developing these kinds of models. Huge numbers of models are being developed. As a medical specialist, I now only use one actual model on patients; only 1 percent of all models developed reaches the clinic. This is because it is very complicated to make a model truly implementable in the clinic. The model must perform well, be certified, be well integrated into the electronic health record, and contribute to better health at lower cost.'

'We need everyone's expertise in every step of the process'

Mark Hoogendoorn

Senior lecturer in Artificial Intelligence, VU University Amsterdam AI research

'Before I start working with my research group to create a predictive model for healthcare, doctors like Paul first try to define for us what we need to look for in the data of the ICU patients. Then we will look at how to apply machine learning [one of the research fields within AI, ed] for that specific problem in healthcare to create a predictive model. We research which model works best or how we can improve

By looking at large amounts of data, it is easier to determine which treatment will work best for an ICU patient.
ANP

technologies in such a way that we can make better models than with existing technologies. It's not surprising that our models aren't used in the hospital just like that – you have to meet strict requirements before you can actually use a model in healthcare, and this also requires many iterations with the medical experts. Of course, this makes perfect sense, but it does make it more difficult to implement a model directly.

Pacmed, Willem's company, tries to get the results one step closer to practical use than my research group. Pacmed actually creates the software for a hospital. Now it may sound like we discover something and Pacmed elaborates on it, but there is a lot of overlap. We need everyone's expertise in every step of the process. It's about cross-pollination. And that's the strength of this collaboration, I think.

Developments within machine learning are needed in order to be able to create new developments. Certain technologies within machine learning are not developed enough yet. We certainly can't solve all the problems yet. For example, some models require an awful lot of data. We want to change this; we want to develop technologies that reduce the amount of data required to get a good model.'

'For the final decision, the doctor remains indispensable'

Willem Herter

Founder of healthcare software company Pacmed, Development of algorithms and software

'A lot of attention is being paid to research into the use of machine learning in healthcare. Machine learning is very promising; with machine learning, computers can recognise patterns in the characteristics of a patient and treatment from large amounts of data and associate them with a possible outcome. On the basis of this, decision support software for the healthcare sector can be developed, which can serve as an additional source of information for the doctor. However, this information must be presented to the doctor in a useful and safe way. Collaboration with the healthcare sector is essential for this. That is why we are very happy with our collaboration with Paul and his colleagues. In addition, Mark and his research group have been an enormous help in the scientific validation of our methods.

Our ICU software is CE certified, and may, therefore, also be used in the clinic. Designing software and documentation that meet the requirements of CE certification is an important step and takes a lot of time and money. Very little other CE-certified machine learning software is available in the ICU. Much of the research done into machine learning doesn't result in practical use, and, ultimately, doesn't reach the patient.'

We believe it is important to present the information in our software as transparently as possible. Doctors, therefore, receive a substantiation of the predictions shown by the software, and, in this way, they can place the information within the entire context of the patient. It always remains the case that the data on which the algorithm is based does not reflect the full situation of the patient. For the final decision, the knowledge and skills of the doctor remain indispensable.' ■