



Interview with Frank van Harmelen on the Future of Generative AI and the Role of Conceptual Modeling

Downloaded from: <https://research.chalmers.se>, 2026-04-15 09:44 UTC

Citation for the original published paper (version of record):

Fill, H., Horkoff, J., Fettke, P. et al (2026). Interview with Frank van Harmelen on the Future of Generative AI and the Role of Conceptual Modeling. *Business and Information Systems Engineering*, 68: 85-88.
<http://dx.doi.org/10.1007/s12599-025-00980-0>

N.B. When citing this work, cite the original published paper.



Interview with Frank van Harmelen on the Future of Generative AI and the Role of Conceptual Modeling

Hans-Georg Fill · Jennifer Horkoff · Peter Fettke · Julius Köpke

Published online: 27 February 2026
© The Author(s) 2026



Frank van Harmelen is a full professor at Vrije Universiteit Amsterdam. He completed his PhD in Edinburgh and then moved back to Amsterdam to work as a postdoc. Since being appointed full professor at the Vrije Universiteit

Amsterdam, he has been researching what is now called symbolic AI or knowledge representation. This involves representing the knowledge humans have and use to solve tasks and solve problems. This knowledge is encoded in a form that is also useful for machines, enabling them to use this expert knowledge to solve problems.

BISE: Thank you very much for joining us for this interview. What is your perspective on the connection between the BISE community, particularly the business informatics community, and conceptual modeling, and your field of study?

Van Harmelen: I think the field of conceptual modeling is a nearest neighbor of knowledge representation. The notion of ontologies for example is a way to structure different types of entities and relationships. There is therefore an overlapping interest between the conceptual modeling field and knowledge representation. I'm a little less familiar with the business informatics field, although some of my colleagues are active there.

BISE: You already mentioned some of the historical developments. So how do you perceive the current progress in generative AI particularly from the historical perspective?

Van Harmelen: I think we should be honest about this, and the explosion of the capabilities of generative AI has really taken all of us by surprise. Nowadays, if journalists ask me to predict what will happen in five or ten years, I tell them I stopped predicting the future in November 2022, when ChatGPT hit the scene. We had talks in Amsterdam by people from OpenAI and they said they were surprised. From a historical perspective, the transformer model and the underlying algorithms were known since 2017. They were in the open literature; people were playing with them. Apart from some small optimizations, the main step was to put in more compute power and more data. And everybody

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

H.-G. Fill
Research Group Digitalization and Information Systems,
University of Fribourg, 1700 Fribourg, Switzerland
e-mail: hans-georg.fill@unifr.ch

J. Horkoff
The University of Gothenburg and Chalmers University of
Technology, Gothenburg, Sweden
e-mail: jennifer.horkoff@gu.se

P. Fettke
German Research Center for Artificial Intelligence (DFKI),
66123 Saarbrücken, Germany
e-mail: Peter.Fettke@dfki.de

P. Fettke
Saarland University, 66123 Saarbrücken, Germany

J. Köpke (✉)
University of Klagenfurt, Klagenfurt, Austria
e-mail: julius.koepke@aau.at

has been very surprised. In historical terms, it has really been a revolution and an event with major consequences that nobody saw coming. But of course, science only progresses in small steps and not really in revolutions, even though people like to think this. The transformer architecture goes way back to the work on deep learning, and the work on deep learning goes way back to the early work on neural networks. So there are deep roots in AI for the current generative AI techniques, but how well they perform has been, from a historical perspective, a surprise.

BISE: Currently there is also a new trend of combining large language models and generative AI with more agentic systems, so that the machines do not only have the possibility to produce speech, but to really perform agentic AI. What is your view on that? Do you think that it is reasonable to combine agentic systems with large language models? How do you see the development?

Van Harmelen: I'm very split on this. Half of me is a scientist and half of me is a citizen. As a scientist, I'm very excited by this. I think if we really want to study computational intelligence, we must study computers that act in the real world. They can be robots that manipulate objects in the real world, or they can be software that manipulates information in the real world, for example, that transfers money, buys items on the web, sends an e-mail, or deletes a file. We cannot just study intelligence by looking at a brain in a computer that is not connected to the outside world. So, as a scientist, I think that this is a very exciting and important development where an AI system can perform actions in the real world and then also learn from the consequences of this action, which is, after all, how humans learn. You do stuff, things go wrong, then you learn from the consequences of your actions. As a citizen, I'm scared because the current technology is so unreliable. It is not just unreliable because it is currently imperfect and we will fix it later. I think it is inherently unreliable because of the probabilistic nature of the algorithms. I think it is irresponsible to allow agents that are so unreliable to perform potentially irreversible actions in the real world. I'm very concerned about the unregulated manner in which the current generative AI, in general, is allowed to roll out, and that concern is even greater when it comes to agentic forms.

BISE: Do you see some killer applications for agentic AI?

Van Harmelen: I think there are so many. I spend my time, even just on my laptop, doing all these boring things, you know, moving files around in folders, organizing things, and changing the date on files. I'm sure an intelligent assistant would help. So, think of it as an intelligent assistant. We could all think of intelligent assistants in our office environment or in our daily life or the simple thing of doing some shopping on the Internet. I make the choice,

and then after having made the choice, I need to do all this other boring stuff: I need to fill in my details, I need to transfer money from my account. Clearly these tasks are different every time, so they're not easy to automate because every website has a different workflow and the buttons always look different. You need some intelligence to do this. You need an intelligent personal assistant. So, I think this personal assistance would definitely be a very attractive use case, but it would require a level of reliability and trust that we currently don't achieve.

BISE: You already mentioned the close relationship between knowledge representation and conceptual modeling. Conceptual modeling has a particular focus on the support of human understanding and communication. How do you perceive the role of conceptual modeling in this context?

Van Harmelen: The single word that comes to mind is *crucial*. It's absolutely essential. Why? We want to be able to collaborate with these intelligent agents. We don't want to just automate stuff and then never communicate with them again. We want to collaborate. And you can only collaborate with an agent if you have a shared conceptual space. You and I can now have a fruitful collaboration, a conversation, because we share a conceptual space. This is the basis for our conversation and collaboration and we constantly work on this conceptual space. We expand it and we debug it and we clarify it. This conceptual space needs to be on the human side, but also on the machine side. I think that the field of conceptual modeling has a lot to contribute to the kind of representations and models that will be needed on the computer side to meaningfully communicate and collaborate with humans. Clearly, in neural network technology, this is completely missing. In a sense, both the strength, but also the weakness of neural networks is that they don't do any conceptual modeling. This is a strength because conceptual modeling is difficult and it's expensive and it's a non-trivial activity, so it's great if we can skip it. The neural network just learns whatever it learns, and we have no idea what it has learned, but it gives the right answers. That's the strength, but it's also the weakness because we have no way of communicating with the neural network apart from the output it gives us, as we don't have a shared conceptual space. I think conceptual modeling needs to be part of this human machine collaboration to build the conceptual space that is shared between the two collaborators.

BISE: You already mentioned the challenges of AI today and some possibilities how conceptual modeling might tackle these challenges. Maybe we could delve a little deeper into that topic.

Van Harmelen: Some of the things we have discussed before will come together. We have already mentioned that the current techniques for generative AI are inherently

unreliable. They are probabilistic predictors, probabilistic algorithms. We all see the hallucinations, the mistakes, etc. I mean, they're great, don't get me wrong. I use them all the time. I have a browser tab on Gemini open permanently. And my university has a company account. I'm not saying that generative AI is useless, but it is also inherently unreliable. We must do something about this, and I think that's a big challenge and that the solution to this challenge will come from the symbolic side of AI.

Which brings us back to conceptual modeling being part of the symbolic part of AI. We want to be able to put constraints on a neural network. There are some answers that will always be wrong, e.g., never prescribe this drug to a pregnant woman, never do that when you're in a self-driving car. Ensuring limits on the behavior is a combination of symbolic and non-symbolic AI. Building the shared conceptual space so that we can communicate with them in terms that are meaningful to us is part of the combination of symbolic AI and non-symbolic data-driven AI. I think the big challenge now is to take these two families of AI: the data-driven statistical neural networks and the knowledge based, logic-informed conceptual modeling symbolic AI, and to somehow combine these two. Some people call this the 3rd wave of AI, with the first wave of AI being the symbolic one, the second wave of AI being the data-driven one, and now the third wave of AI trying to combine these two. It's technically very difficult because they use very different representations and very different techniques. Even the underlying mathematics is very different. But I think that will be a very exciting challenge.

BISE: Do you see a particular influence when you focus on your own research, knowledge representation, or conceptual modeling? Do you think GenAI offers important improvements to this type of research?

Van Harmelen: Certainly, as I said earlier, we are now no longer doing pure knowledge representation and conceptual modeling. We are combining it with machine learning in two directions: we are using knowledge representation and conceptual models to improve machine learning and the other way around – we're using machine learning to improve knowledge representation. Let me give an example of both. Machine learning algorithms use insane amounts of data, and I am always very annoyed when I hear a colleague from machine learning tells a doctor "Oh, no, sorry, I cannot solve your medical problem because you don't have enough data." I think, wait a minute, you have the wrong algorithm. Humans can learn much faster. We don't need to see a million pictures of a zebra before we recognize a zebra. This is because we use background knowledge to inform our learning, so we can really inject background knowledge into machine learning algorithms to use less data to learn. That would be an

example in one direction. An example in the other direction is how could we use machine learning to improve conceptual modeling? Well, conceptual modeling is hard, it's expensive. People build conceptual models based on text and interviews with experts. I think generative AI can help extract a version zero of a conceptual model from a corpus of text in order to make the process of conceptual modeling faster, cheaper, and more accessible to non-experts. Machine learning can help conceptual modeling; I think the interaction is really in both directions.

BISE: We are increasingly seeing different sectors converge. If you look at today's large language models – mostly developed by major US and other international tech companies – you can see that, in terms of how they develop and disseminate knowledge, they are moving into what has traditionally been the domain of universities and academic institutions. Universities have long been responsible for sharing knowledge, teaching it, and advancing it through research. Yet, these large language models are not products of universities; they are created by private industry. In this context, how do you think universities and academic institutions should position themselves? What role ought they to play in this evolving landscape?

Van Harmelen: I'm very happy that you mentioned teaching in your question. I think universities are still figuring out how to use these tools in teaching. Clearly, we don't want our students to hand in an essay written by ChatGPT. We need to teach our students that the essay is not the point. The real value of writing the essay is the process that they have to go through to write the essay, and then the essay is a kind of a byproduct. The teacher reads the essay and then nobody ever sees it again. The learning process in the head of the student is what really matters. We have to teach our students how to use these techniques wisely because they will use these techniques in their future professional lives, but we have to also make sure that they don't use them in a way that destroys the teaching process, and I don't think we have the answers to that.

I think it is also problematic that all the innovations now come from a very small number of companies, and they are profit driven. There's nothing immoral about this, that's why we have companies, that's how we organized our society. But they are not necessarily serving the public good. The domination of a few companies is partly because the current generation of generative AI requires massive infrastructure, as well as huge amounts of computation and data. Universities are simply not capital intensive enough to ever compete with that, I fear.

BISE: Recently, some Swiss universities provided the Apertus model as the first fully open source LLM whose training data and methods are disclosed. It is not yet on the

level of the commercial models but constantly updated. What do you think of this direction?

Van Harmelen: I think that's great, and I really admire Switzerland for doing this. In the Netherlands, we are also trying to do this for the Dutch language: GPT-NL is trained on a legally acquired text corpus in Dutch. I also hope that the current phase of these very large models is a temporary phase. Maybe in five years, we'll laugh at how we thought we needed to fill Arizona with server farms. There has been significant progress in developing smaller language models. Using different techniques, smaller models can be quite competitive. Additionally, there are multi-agent models, which are a mixture of experts. Why would we need a single model that is good at everything? Why could we not train one model that's good in medicine and another one that's good in computer science? Why would a single model have to be good at both? If we move towards smaller models and this mixture of experts, then maybe we can reduce our current dependence on very large capital-intensive infrastructures and this would create a healthier balance between academic research and commercial interests. On my optimistic days this is what I am hoping for. Otherwise, we must follow Switzerland's lead and try to build our own infrastructure.

BISE: As established researchers we all have ideas what are the grand challenges and opportunities of the field. But if we particularly focus on young researchers, do you have some advice? What are the grand challenges and opportunities? What should young researchers from the conceptual modeling or knowledge representation community focus on?

Van Harmelen: That's not an easy question. Sometimes, I sense a tendency in some parts of our community to resist some of the new developments. Then you end up in these GPT-bashing conversations: "they cannot do this, and they cannot do that, I asked them this morning and they gave a really stupid answer". I don't think that's a good attitude. These generative AI techniques, they are with us. They are not going to go away. They're too good. They have really changed our understanding of computational intelligence. So advice number one, don't ignore them. In some corners of the knowledge representation world, there are good papers, but there are also papers that could have

been written five years ago. They are not bad papers, but they ignore the changes in the world. Don't ignore the changes in the last few years. The other advice would be: don't leave symbolic AI, because I really think that symbolic AI will be part of the next improvement steps in computational intelligence, for the reasons that we discussed before: explainability and safety and learning from fewer data points. If you look at the master courses at some of the leading universities, symbolic AI has just been completely removed from the curriculum. I did my PhD in Edinburgh and my Alma mater has almost no symbolic AI in the master's program anymore. In some universities, also in the Netherlands, this is the case. If you are in conceptual modeling, if you are in knowledge representation and symbolic AI you have a rare skill and you should keep that skill because it will be needed, but it will be needed in combination with these new techniques.

BISE: So you strongly vote for a combination of non-symbolic and traditional explicit modeling approach?

Van Harmelen: Absolutely, and conceptual modeling is an integral part of that, yes.

BISE: Thank you very much for the very interesting interview.

Funding Open access funding provided by University of Klagenfurt.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.