



AI Ignition

Ignite your AI curiosity with Kristian Kersting

Beena Ammanath: Hi, everyone. In this season of *AI Ignition*, you will hear new voices and insights from my Deloitte colleagues. Now, I will turn it over to our guest host today. Enjoy the show.

Björn Bringmann: Thanks, Beena. And hi, my name is Björn Bringmann. I'm the managing director of the Deloitte AI Institute, Germany. Guest hosting *AI Ignition*. Today on *AI Ignition*, we're joined by Dr. Kristian Kersting. Kristian is a Full Professor at the Computer Science Department of the Technical University Darmstadt in Germany. His main research interests are statistical relational artificial intelligence, as well as deep and deep probabilistic programming and learning.

In addition, Kristian is the head of the Artificial Intelligence and Machine Learning [AIML] lab, a member of the Centre for Cognitive Science, the faculty of the ELLIS Unit Darmstadt, and the founding co-director of the Hessian Center for Artificial Intelligence (hessian. AI). Beyond this work, Kristian has published over 180 peer reviewed technical papers and offered and edited multiple books on AI and machine learning.

Welcome, Kristian.

Kristian Kersting: Thanks. Thanks for having me.

Björn Bringmann: Kristian, today you are a thought leader in AI machine learning pretty much across Germany, and I would claim beyond. And I do know you for quite a while but, in fact, never asked what was the catalyst that brought you towards that career? How did it start?

Kristian Kersting: There was one moment where I really felt very attached to computer science, and it is, you know, my mom was not happy with me. My room was very crowded again. And so she told me, please clean finally. And I was wondering, why do I have to do it? Why can't there be a machine? So I was wondering whether AI could be the answer to give some form of order to chaos.

And I think that is a question I still have. I don't know whether I can give you an answer, but that was one part there. So how can we optimize cleaning? What is intelligence? But to be honest, I never expected to become a professor, and maybe at some point I had no other chance anymore because I felt like maybe I'm not good enough for real business.

Who knows? But that was one of the reasons.

Björn Bringmann: Wow, OK! An unexpected story, for sure. And the Euro professors are known to have quite a grip on chaos, right. So, that probably fits. So nowadays, I checked, your wiki page yesterday, but also on your... your team is working on pretty much making computers learn like humans do, right? At least, that's what I understood. So making them learn as rapidly as we do. And, on the other hand, as flexible.

Kristian Kersting: Yes.

Björn Bringmann: I remember you've done some things in the past, but how did you get to that one?

Kristian Kersting: Well, the flexibility, I think that's what we are currently all experiencing, right? I mean, all these large language models, transformer-based approaches, but also diffusion models are so impressive in the flexibility they can be used and they also show. I don't think we are already at the stage where they are as efficient as humans. So now, why do I care about that?

Well, AI has a twin discipline—and it's cognitive science. So cognitive science is essentially asking, can we find AI rhythms that explain human behavior, smart behavior, intelligent behavior. And so we can learn as AI, we can learn a lot from humans and how humans are doing it. We may not have to always follow the human way, right? I mean, smart behavior can be completely different.

And if you think about it, that's also what you want. Because if you think of self-driving cars, if you're just with the human way, then we have as much as much [road] congestion as before. So it's not about being super human. It's about just solving issues we have in a better way. Think of many, many things in climate change.

How can we do that without having the human bias maybe in there? Can we get supported at least? Can we have a system that can be an expert in many domains and help us solving all our problems? Maybe we can ask the system. So this is where the human dimension in the sense of cognitive science is getting in there.

And maybe we need heuristics, which are maybe not actually heuristics to deal with the complexity of the real world, even within machines. To make it more concrete, if you think about how much a single query to a large language model currently costs, it's 10 times more energy costly than just doing a web search. If I ask you, I think, maybe you start [to] sweat, but you're still much more energy efficient than any of these systems.

And also, you don't need millions of examples to understand what is a cat and what is a dog. So there seems to be still something doable here to make our machines more efficient. So this is the driving force, and we go step by step there, right? How to get reasoning and learning together or think of playing Atari games.

We are so good in AI using deep learning to play Atari, but then they are all pixel-based. And if you think of *Breakout*, it's the system where you know you have this little paddle and then the little ball and you have to kill the bricks in the upper part. And if you then move the paddle a bit up, most of these deep learning systems break because they never experienced that situation because they learned that the paddle is somewhere else.

So how can we make it more object-centric? And for us, it's so natural: If we talk to each other, we are using the notion of objects and even relations. So how can we get that together with deep learning together with pixel-based? How can we make these systems more robust? This is what we mean by make it more... learn more rapidly as humans and also more flexible. And finally, the big, big questions we are currently touching is can these systems understand our social norms so that they are not confused by our behavior? And I think this is pretty cool—and also because we learn a lot about our own behavior. So social norms is something we are looking at.

Björn Bringmann: Wow, that was a lot and quite some things to unpack. I took some notes, so I don't forget anything. But if I understood correctly, you would see AI—I'll just call them "tools" for now—as complementary, right? I mean, you said they don't need to do the same as we do. So there's this other question in there obviously on the current progress and big discussions on if they outperform humans or not. Do you have maybe a short statement on that?

Kristian Kersting: Yeah, to make it very short, yesterday was the first birthday of searchability. It's one year old and still it can't—is that the correct English? The right English? It can't "blow off the candle" on the birthday cake, right? It can't simply turn it off. There's no way because it can't directly interact on its own with the physical world.

So there seems to be still a lot missing on AGI [artificial general intelligence], but it's impressive what these systems are capable of doing and how they can help us solving our issues. But we still solve, in the end, the issues. And this is mind blowing. I mean, I expect that there's no language barriers anymore. I expect that we can get access to all the different scientific disciplines in an interactive way.

But still we have also to learn a lot ourselves about the other disciplines. So it's amazing. But yeah, and it can't deal with the birthday cake. It can't bake a pizza. There are many, many things these systems can't do currently, so still a lot to be done.

Björn Bringmann: Any thoughts on the Generative AI will grace, so to speak? Like any recommendations? Maybe what to—how to—like, from your point of view? Like, what's the hard part, so to speak, in the hype? Like, what's the best stuff that's just fluff at the moment?

Kristian Kersting: So first of all, I think Generative AI is mind-blowing, and I never expected that we would make a jump in capabilities and performance, and this is why I guess some people talk about AGI because it's so rapid. And maybe there is just a little bit of a step, and then we have AGI. I don't think we have AGI soon, but this still means we can have many, many quantum leaps and enjoy the progress there.

And so I think we should embrace it. But we should embrace as always with a bit of consciousness or with a bit of soul in a sense there, right? So we have to be a bit careful. So, for example, what we are working on is, it's amazing what type of images you can generate by now. And next to social implications, you know, are there still maybe designers needed? Which I think, yes, they are. But still every one of us becomes a bit more of a designer, which is amazing. We empower people—our AI can help to empower people. But we have to be careful because you may also generate a lot of "crap," to be honest. So you may deliver a tool into the rooms, into the homes, to generate... harmful content, and you don't want that.

And we want to have systems that then tell the user, "No, I will not produce it because of." And gives a justification of why not to use. And I think that is also important for all the issues maybe related to misinformation in the upcoming elections worldwide. I think, next year [2024], we would have something like 2 billion people on elections worldwide.

And so we have to make sure in the long run that these elections are not heavily influenced by Generative AI. So that works for the image part. For the text part, similar stuff, right? We have a system published at Nature Machine Intelligence that tries to reduce the crap that is produced, but there still to be done a lot. So we were just showing it's possible, we believe it's possible, but still a lot to be done.

And also in terms of providing justification, links, combining it with our reasoning—so going a bit into *can't*. You want to have a system that may abstain and say, "No, I don't want to do that," but not on an emotional level, right? What we call in Germany *genug*, right? This is... you want to have reasonable AI. That's what we are working on.

Also, it's amazing what you can do. Now, in terms of business, I think people have to be a bit careful. That's why I was talking about AGI. I think they should not put AI systems on the same level as humans already. This will take ages still to come. So if you don't do that, that would be a good step because then you can enjoy what we can do already.

Björn Bringmann: So you're saying use them for what they are, not for what you want them to be, right?

Kristian Kersting: Yes, yes. I think that's really—we should not call for short, "Oh, yeah. AI equals humans." That's the wrong picture.

Björn Bringmann: Yeah. And it supports, so to speak, the idea of having some complementary system, rather than rebuilding a human, but just with other means.

Kristian Kersting: Yeah, exactly. There's the slogan, "From AI to IA." So from artificial intelligence to intelligence amplification. And I think that is what we should think about—how to get—if we want the robot out of our daily business, how can we get that onto a robot so that we can spend our time—if you're a medical doctor, for example—really with a patient or with doing even more research, and less on filling out forms. Or if you think about text forms, why do we have to fill them [out]? I think they can be done automatically with some information that we want to keep private. Or if you think of the admin stuff—at least at the university, at a German university, you're spending a lot time on administration.

And I think this is where the systems can be very useful to prefill. And in the end, I'm still in charge to say, yes, that's the right form and it's correctly filled. That would be nice. Next steps. And that's not—I mean, you may say it's super, super human, right? [laughs] Because if you said building forms, I think you're super human! But it's not this "super human" where people feel like we want to find the next level. I, at least, I'm just impressed already by this help that an AI system can do.

Björn Bringmann: Is there anything we would say should prepare to... It's actually two questions in one here: one, a little bit of an outlook, but maybe let's pause it for a moment; and the other one, regulations from our government. Like, everybody is called up on the radar at the moment, so to speak. So anything you would say we should prepare for—business should prepare for—to take this into account?

Kristian Kersting: I think regulation will come, and it's very natural. I don't understand why we should be scared by regulation per se. That's what we're used to. In medicine, in the medical domain, in physics, and in so many domains, we are used to regulation, and that's fine.

So instead of just talking about regulation, I think we should also talk about how to set up the production line. And I think we need open models. And then from there we develop, in particular, closed models that can then be used by companies. I think that's more important. But regulation procedure? I think it's very natural. I don't know whether I would regulate the base models, foundation models. I don't know. I'm happy already that we have the application-oriented one.

More important is that we need open models. It's not only me as a more technical person who should give an answer; it's a legal question, it's a philosophical question, and many, many different voices in there. For me, I rather think we should go ahead—and I feel like we compare again with humans at many levels here already. We should just make sure that it's testable to some extent, that we have it open, that we can ask the model things even if we are not the company. That is more important for me than saying it has to be fully regulated, and all the models have to be completely open, because, you know, I look at my lamp here on the desk and I don't know whether I fully understand how it works, but I'm using it. I'm trusting it, right? So it's much more about trust than being completely transparent, and we have all details. That doesn't help me.

Björn Bringmann: Yeah, makes sense. I also don't really understand all the details of my car, but I trust it pretty much. And I do have a rough understanding you're also serving as a school director of Hessian AI, right? I was just wondering if many of the reasons you just listed are part of the reasons why—I think you even co-founded it, right? Or at least part of Hessian AI. Is it related to what you just said? Making it more public and not necessarily controlled, but probably also educate people? I don't know.

Kristian Kersting: Yeah. So I strongly believe that AI needs an ecosystem. An ecosystem needs also different voices. And then we proposed to go for a larger AI center, and luckily we got money for that. But I have to say it's also pretty tough because there are 13 different universities involved. It feels a bit like only what was left on the line currently has a more horrible job in terms of, you know, keeping different opinions together and negotiating them and coordinating them. But I think it's use is useful. It's really, it's a good effort that you spent there. It's something good that it's getting out of it. Because you get more a bit more robust, you're getting more responsible in a sense, and you understand that there are different opinions, and you're not just going with your own opinion. So, that was one part. But then AI per se needs an ecosystem also in terms of unique academia, but you also want to monetize. So you may need entrepreneurship, so we also have entrepreneurship involved. But then all that doesn't help you if you don't have the infrastructure to compute something.

So we are also happy that this year we got *fortytwo*, which is probably the largest academic AI supercomputer, at least in Germany. Not sure about... maybe France has more. I don't know. That's a different question, but it's really just for AI. So we are very happy there. We have also produced already a German language model, and using that, there is some other great news coming.

I'm just happy to get different opinions involved because that makes it more robust and much stronger, and many questions in AI need an ecosystem. The questions are so big that you cannot say it's a single professorship or a single company if the company itself is not very big. That doesn't work.

So we need all these different teams that help each other, that come together for a particular problem and then maybe form another team for another problem. And that's what we are working on. Infrastructure was a big step. I'm so happy about that. I find it very useful.

Björn Bringmann: Is there anything from your interactions with different companies and so on that you've learned what you would say there's the opportunity. Or that's a good approach on how to start. I know not all but many companies are still kind of starting, right? They may already do quite a lot of stuff, but few have really figured it out. So anything from what you've seen on how to get started, what you would—maybe also why to get started?

Kristian Kersting: Yeah. So I think the easiest thing I can say—and it's maybe not that easy—you need a slightly different culture, and let me explain it. Maybe the best way is to tell you a story how not to do it, and it's a real story.

We were at a larger meeting with many, many, even really big companies. I will not say any names. And one of the—I don't know what—a CEO or CTO, I don't know. But anyhow, he was saying, "OK, problem problem, problem, problem, problem, problem, data, data, data, data, data, data, algorithm, algorithm, algorithm, algorithm, algorithm. You tell me for which problem, which data, which algorithm."

And then all the AI people in the room were a bit like, "Hmm, hmm, eh, and no," almost starting to laugh, and then saying no, sorry, this is why you have to go to university to some extent, or at least to spend some time in studying AI because we can't map every problem to a particular type of data and a definite algorithm to use.

You can actually prove that there is not a single best algorithm, and you need to develop a smell in a sense and understanding for what are maybe good candidates. And then he said, "Ohhh, OK, I understand. So let me repeat: problem problem, problem, problem, data, data, data, data, algorithm, algorithm, algorithm. You tell me blah, blah, blah."

So I think this is what we have to avoid. We have to accept that computer science and AI is not just nerds. It's a science. That's why it's called computer science in English, right? And not just in phonetic, but sometimes you call it... It's a science and something equivalent to mathematics, to medicine, to physics. So please just treat it as a cool science.

Next step is change culture. Don't ask "How can I solve my issue in the traditional way using AI?" Just ask "Here's an issue. How would you solve it?" Because then you get the freedom to break out of your processes and maybe come up with a new process which is AI-driven.

So just go out, give it a try, but don't expect within two weeks "Now everything is different." It takes time, it takes effort, it takes money. But that's fine.

Björn Bringmann: Very, very cool. That's super, super helpful. Obviously a bit of an outlook. What would be your mid-term, long-term—you pick your distance—predictions for AI and machine learning?

Kristian Kersting: We will have super useful, powerful AI systems that are really changing the world. There will be no language barrier within the next, I think, even maybe three or five years, maybe more like three years. I think really, like Enterprise, you know, the [Star Trek] movie and series, I think you will have this translator automatically and if you go almost anywhere on Earth. We have to make sure to not forget many of the languages, right? I mean, it's not just English, and it's not just German, and it's not just French and Spanish. There are many—in Mandarin. There are many other languages we want to cover, but I think there will be no language barrier for any major language soonish, which will be amazing.

I think in other stuff that will happen, we will have a system that fills out forms for you. Really. I really believe admin will be much easier for many, many, many people. And that also includes—and that's maybe very interesting for Deloitte—you will have systems that help you to come up with much better presentations. Maybe they come up even with a presentation, and you see already the first tools doing that, and then this will all happen very soon.

Ten to 20 years, I think we will see improvements—maybe a bit more like 10 [years]—a lot [of] improvements in robotics, but it's not as easy. I think it will be much harder for that. We will see systems that can do reasoning. But really with guarantees and learning at the same time, and the interaction will just be by our voice. But they will also understand our mimic. They will also understand our reactions. And so, this is still a bit [of] work.

What will not work within the next 10 years, I think, is a machine that can tell jokes well and understand jokes. So the understanding jokes is more important, right?

Imagine that we all go to Mars, which I don't believe we should maybe do. I hope there will be other planets because it's pretty hot on Mars. I think it would be a boring planet. But anyhow, if you have to travel there, a system that is monitoring the crew should also understand, you know, how is the atmosphere? Do people like each other? What's going on there? And that means there has to be empathy, there has to be understanding of emotions. And that's pretty tough because do we want machines to understand it, to maybe even trick us? I don't know. But that's very hard.

Björn Bringmann: And actually what you just mentioned at the end kind of nicely closes to where you were at the beginning, right? You said you're working quite strong or quite intensely on having machines or enabling machines to understand social norms.

Kristian Kersting: Yeah, yeah. And it's also being aware that it's super tough, right? It's not—I don't want to say it's easy or solvable, to some extent, soonish. We just try our best, and interdisciplinary research is required there.

Björn Bringmann: Mm-hmm. For our listeners here, how could they get connected with you? Where do they find you?

Kristian Kersting: They can find me on X formerly known as Twitter. They can also send an email to my secretary or to @hessian.AI. We have something like info@hessian.AI, I guess. They can send emails there. They can also maybe ask you, and for very prominent people I'm happy if you even forward the handy smartphone number.

There are different ways, but that's the easiest way. Just drop an email. If you are close to Darmstadt, drop by—super happy to talk to people. Or, Björn, we could try to organize a symposium together or something like that, right? I mean, there's maybe also something to be done there.

Björn Bringmann: Yeah, if somebody is interested, reach out and then we'll see how to set this up. Pretty cool. Kristian, thanks again for being with us on the show today. And thanks to our audience for tuning in to *AI Ignition*. Be sure to stay connected with the Institute for more Research and Insights.

Visit the AI Ignition Podcast Episode Library
[Deloitte.com/us/ai-ignition](https://deloitte.com/us/ai-ignition)

About Deloitte

As used in this podcast, "Deloitte" means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.

Copyright © 2024 Deloitte Development LLC. All rights reserved.