

**KIMBERLY** Welcome to *Pondering AI*, where we contemplate what is required to ensure artificial intelligence is deployed fairly, safely, and justly for all, now and in the future. I am your host, Kimberly Nevala, and today, I'm so pleased to be joined by Teemu Roos to explore the path to pervasive AI literacy.

Teemu is the lead instructor of the wildly popular "Elements of AI" online course, and leads the AI education program at the Finnish Center for AI, all of which is in addition to being a professor of computer science at the University of Helsinki, where his research focuses on future applications of machine learning. Hyvää päivää, Teemu.

**TEEMU ROOS:** Hyvää päivää, Kimberly. That's brilliant.

**KIMBERLY** --with apologies for my poor Finnish pronunciation. So tell us how you came to study computer science, data science, and artificial intelligence in particular.

**TEEMU ROOS:** OK, so how long do you have? It's a long story of course, but--

**KIMBERLY** (Laughing) Maybe just the highlights.

**NEVALA:**

**TEEMU ROOS:** Right. So I was always really curious about computers. It wasn't really the gadgets. It was more like the creativity, that you can create these imaginary worlds by programming and making things happen. It was, kind of, like magic, as a kid, to make something happen like that, and that was always the thing that drove me to learn more about it.

And I went to university to study computer science, and I was on my path to becoming a programmer, a developer. And then, during my studies, I took this part-time job as a software developer, but that was, kind of, a turning point when it turned out that the routine work of a software developer can actually become pretty damn dull. At least, I wasn't able to create new things, because it was always specified to the point that it was a billing system for this large teleoperator.

So there was not much creativity, as you can imagine. And I actually noticed that I enjoyed my studies, and I was studying-- we had some courses on AI, Artificial Intelligence, and machine learning, that really caught my attention, and I really got taken away by that. And then, I was thinking, OK, well, maybe the software developer work isn't the coolest thing or the funnest thing in the world, but I could maybe continue studying a bit longer.

And then, I became a PhD student just to enable myself to learn a bit more and so on. And that's basically what I'm still doing, and I'm still learning more as a researcher, as a professor. It's, kind of, a dream job, because I can always continue that forever, always learning things that maybe somebody else hasn't learned before, so creating new knowledge. And that's the long and short of it.

**KIMBERLY** So it sounds like you've got the best of both worlds. In fact, your actual job is to learn and to help the rest of us learn.

**NEVALA:**

**TEEMU ROOS:** Exactly, exactly.

**KIMBERLY NEVALA:** So you're doing some really exciting work in two very divergent-- maybe not really divergent areas-- but in both applied machine-learning research and what I'm going to broadly label public education. Can you talk to us first about what are some of the projects on the research side that you're most excited about today?

**TEEMU ROOS:** Well, OK, so today, literally today, we had this brilliant meeting planning joint research with astrophysicists. And you know, it's quite exciting, even to me. Of course, I'm sort of a layman in terms of physics or especially astrophysics, so all this kind of magnetohydrodynamics sort of-- yada, yada, things-- that's, kind of, magic still to me.

It's less-known magic than the AI magic. And now, it's the opportunities to combine my machine-learning research and applications that are quite exciting. There's another project where we're looking into neuroscience where really small babies' brain development is being monitored by connectivity patterns that are analyzed with machine-learning technologies.

It's really the best part in addition to being really fascinating in itself. But another best part in machine learning is that you can apply it in almost everything-- almost in actually everything. There's no field of research where you couldn't find something where you could help with machine learning. That is really inspiring.

**KIMBERLY NEVALA:** The potential, I think, for all of those things is so outstanding, and it struck me as you were talking about babies, that we're using machine learning to learn how we learn, to some extent. So that observation that machine learning and artificial intelligence can really be used and intersect with all aspects of life and all careers and all of these different paths. Is that what inspired you to bridge the gap between the work you're doing in the research lab and commercializing AI and more broad efforts for literacy and education?

**TEEMU ROOS:** Yeah. Yeah, yeah, so that's quite right. And there's of course a lot of work in applying machine learning in physics or engineering. But I was always feeling that there is so much lost potential, because the community is not quite as diverse as it could be. So there's really heavy focus and a lot of investment of effort into certain kinds of applications-- exactly those gadgets.

As I said, I'm not so keen on the latest gadgets. There are other things that appeal to me in computer science. And I've felt that it's kind of a little bit too homogeneous in terms of the interests of the people that are involved.

And that's been something that I've been wondering about for a long time. How do we get more diverse people, more diverse community, in creating those ideas? Because it's not as if I, as a machine learning person or a computer scientist, could come in a new problem domain-- let's say, medicine or design or teaching-- and start creating viable and useful ideas there.

It always takes the person who is immersed in that topic and has been facing the problems in that topic and in that domain for years and years to see what is missing. And that's why it's always collaboration. As I said, we were just chatting with astrophysicists, and I couldn't dare pick a book on astrophysics, read that, and pretend that I'd be able to come up with the best ideas.

It's always the people that are thinking 24/7, every waking moment at least, of that problem. And they come up with one part of the puzzle, and maybe I, as a computer scientist, then can bring the other part, but it can't be done without them. And I think that meeting of people, meeting of minds of different backgrounds, is when the magic really happens.

**KIMBERLY**

**NEVALA:**

So when you were then thinking about putting together "Elements of AI," which is wildly popular-- I believe that over 650,000 people have now taken that course. I think it was also the most highly rated computer science course, in terms of MOOCs - Massive Open Online Courses - by Class Central. What was your objective, and who was the target audience for that course?

**TEEMU ROOS:**

That's a great question. The target audience-- very briefly-- is everyone who is not already interested in AI. So you would think that's a really weird way to define the target audience for an AI course, people who are not interested in it.

Maybe you'd think, OK, well, it should be the other way around. But we actually thought about it for a long time. When we were starting this collaboration with-- there's a company called Reaktor, who is our partner in the "Elements of AI" project, and we were thinking together with them, who should we have as our target audience? And we interviewed a lot of people.

We asked them why they're interested in AI, why they're not interested in it, what aspects of value they might be interested in. And through those discussions, we realized that there are a lot of people who don't realize that AI is relevant to them. But once they realized that, they actually said, OK, well, I should actually probably learn a bit about it and form an informed opinion and do something about it, be more active about it.

And we found that pattern over and over again, that people didn't really realize. They didn't know what AI means, so they couldn't be interested in it. But then, of course, everybody's interested in what happens to their data when they use social media or when they do a search on the internet, when they use various apps on their mobile phones.

So people are actually interested in the topic itself, but they just couldn't connect the dots-- AI having something to do with those things. And that's why we picked that as the target audience, the people who don't maybe yet see the relevance of AI in their life, but who will be interested in it once they see the relevance.

**KIMBERLY**

**NEVALA:**

So how do you get people who might say, OK, this is the "Elements of AI," and maybe I don't even know what that really means and I'm not sure it's relevant to me to click in?

**TEEMU ROOS:**

Well, that's the difficult bit. If we just let it sit somewhere online, people will actually not find it. That's the question, and that's the issue, and that's why we need to do so much more than create an online course and just have it, sort of, sitting there.

We need to create this initiative. We need to think in terms of the growth of the community, word of mouth. How do we encourage people to do it together maybe with their colleagues or their loved ones, their family, encourage people to join maybe reading clubs or meetings?

How do we encourage different stakeholders to organize meetings around it? So it's really so much more than an online course. It's an entire initiative that we were trying to orchestrate from the HQ in Helsinki with local partners in every country where we've launched it.

And we've got amazing partners that generate those ideas. You know, let's here organize something in the local libraries. Let's here organize a hackathon. Let's organize something with the schools, including maybe in some high school curricula or stuff like that. And really, it's another business to find the people that might be interested in it.

**KIMBERLY NEVALA:** So the course is called "Elements of AI." What are the essential elements or topics that you cover in the course?

**TEEMU ROOS:** We, kind of, start with the observation that it's really hard to define AI in a crisp manner. So it's really hard to say that on this side of the border, this is it, and this isn't. So we're rather linking it to something that might be more familiar to people, like statistics or computer science or software in general, and say that it's a subset of that.

And we characterize it like, well, it's the kind of software that might be adaptive or might be able to complete some tasks autonomously rather than being explicitly programmed to do something. And once we've gotten that a bit more clear, we start pointing at concrete examples like, as I mentioned, social media, of course different robotics applications. But we don't, for instance, start with the robotics, because that would give you the impression that is something really far in the future, or maybe more like science fiction things.

We try to intentionally make it more-- I sometimes say, make it more boring. So make AI more boring. Make AI boring again. That's the goal. Well, not exactly, but still-- more commonplace, more familiar.

And then, we go about a little bit peeking under the hood, saying, OK, what might make it work for this particular application? And we give simple examples of the types of principles that are applied in building systems. And then, finally, we circle back to the social and societal implications.

What does it mean now that these systems are developed and applied? What does it mean, let's say, for the working life? Is the job market going to be disrupted?

What does it mean to our privacy? How could you protect your own privacy? And so it's linking those to bigger contexts that people also can relate to more easily.

**KIMBERLY NEVALA:** Yeah. I'm always-- I don't know if shocked is the word-- but taken aback by how little even some of my family and friends are aware of how their information is used or even that these systems are in play today. And so as a very simple example, a dear cousin of mine, probably just a year or two ago, we were talking, and she was really irate about the fact that she had decided to Google herself (which I thought was funny to start with, because she's not usually online) and was really irate that all of the information on her family tree essentially was online.

So she could see her siblings and who they were married to, and she was really appalled. And I don't know if I'm a good friend or a bad friend, because I didn't want to break it to her that was really just the tip of the iceberg. So as you're going through this course, I imagine people are getting really excited about the potential. Do you also find that people are somewhat amazed or taken aback by the extent to which some of these systems are already in play and how their information is being used?

**TEEMU ROOS:** Yeah. And I think that's part of the plan, to some extent. The kind of impact that we'd like to reach is not really on the axis of good or bad, or whether you like AI or whether you dislike AI. We don't want to start saying that, OK, you should like AI. Or we don't like to say, of course, you shouldn't like AI.

But we like to tell people that they should be more active as opposed to being more passive, in terms of forming an opinion, being vocal about it, talking about it to others, and also, of course, at some point, making decisions about using certain products or not using them, and eventually through the democratic decision-making process, think about, how do I behave in voting or otherwise taking part in democratic decision-making? How do we steer AI to a direction that we like? But of course, we leave it to people to judge for themselves.

It's not as if we could tell them, these are the right values, and you should value this thing over that thing. It's really up to them. But we just like to open up the discussion and clarify and make it easy for them to form this informed opinion.

So it's, as I said, like, you don't want to go telling people that was a right thing to do or that wasn't the right thing to do. Just explain to them what are the potential consequences of, let's say, sharing your data. And then, it's up to them. Some people don't mind, and some people rather do.

And of course, it's also not only your own data. So I think that is one of the things to realize, that if you're sharing something, if you're sharing family photos, it's not, of course, only your own privacy that you're perhaps exposing to violations, but other people's that might be featured in this photo. So just an example of things that it's at least good to be aware of them.

**KIMBERLY  
NEVALA:**

So we're trying to make, quote unquote, informed consent truly informed, in a sense. And I want to touch a little bit on your mention of value. And certainly, the conversation of ethics in AI and what should and shouldn't we apply this to is a hot topic at the moment. But I had made the assumption when I had first heard you speak about "Elements of AI" and the course that it was primarily of interest or taken by folks that were coming from outside of technical fields or folks that were already interested in maths and science.

And in looking at some of the reviews when I was researching for the show, that in fact was not the case. I was really happy and surprised to see that there's a lot of feedback from folks who said: I work in the field and really learned some things and have still found this course very helpful. What are the gaps in the landscape today, in terms of how we teach and educate about AI, that this course fills for folks that may in fact have a more technical background already?

**TEEMU ROOS:**

Yeah. That's really an interesting question, and it puts a mirror in front of myself. Because of course, as a professor of computer science and data science, I'm one of those people who educate some of the professionals. And that's indeed been a surprise and given me a lot of food for thought about how we should educate people that-- I've also heard these stories from people that have, let's say, even an academic training in computer science.

They'd say, I've never thought about these ethical concerns. I've never thought about the links between, let's say, political polarization and social media recommender systems. Yeah.

That is something that I've been thinking now more recently and thinking, well, should I myself do something about it? Of course, and I've started thinking, OK, we should probably highlight these issues more in the computer science training that we give to the students at the university. I myself didn't really encounter issues like societal justice issues in my computer science training.

I did take some courses on moral philosophy or ethics where I would be exposed to them, but at that time when I was studying, I wasn't able to link them really. This is something that I've, in a way, learned along the way as we were building the "Elements of AI." And since then, it has accentuated the need for maybe changing and revising the computer science curriculum to cover these aspects much better than they are currently.

**KIMBERLY**

**NEVALA:**

I think that gap probably doesn't just exist in the computer science curricula. I actually am a chemical engineer by training and happened to apply for a job with a friend of mine at a consulting company and got the job. She didn't. We are still friends, which is nice.

But I don't know that I really thought about things, even in the context of basic risk management, until I was actually out being trained in doing the work. For instance, as a management consultant and doing some of that business strategy work. So it seems to come in pockets. But what artificial intelligence is bringing to the fore is that this is something we need to think about more holistically, whether it's across all fields of study, but I think also in how we approach this within organizations and companies.

So I'm interested in how some of the lessons you've learned and your experiences with bringing people into the AI fold translate into the corporate world and into organizations. I think we've gotten beyond this belief that we can just go out and hire a few AI unicorns who will do it all, the data science, the data engineering, and be able to reflect the business case and subject matter expertise in all of these components. And that we need to have more collaboration. But we still hear a lot about this looming skills shortage and the need to go out - and I think this is particularly troubling for small- to medium-sized businesses - and try to hire these very expensive and still relatively rare skills. Could we be doing more to find resources, bring them into the AI fold, and train and reskill from within our existing workforce today?

**TEEMU ROOS:**

Definitely. And I think the domain experts are definitely needed. It's not as if computer scientists can go there and automate them out of the picture. They'll still be needed in creating the ideas and appreciating the better or the worse solutions, and telling which way to go.

So I would even think that training people that are the experts in whatever field they are experts in is going to give you-- in a way, achieve that more diverse community of people who contribute to the AI knowledge and the solutions that we have. So in that sense, that's exactly what I'd like to see. And I don't think people even necessarily have to learn programming. It's good if you get some exposure, understanding how long it takes to write a certain amount of code, and of course, how much a certain functionality-- how much code that would require.

I think those kinds of things on that very high level is good for a manager who needs to be able to converse and collaborate with the programmers. So I think there's a lot that we can gain if people in managerial positions or other experts who collaborate with the technical experts, the machine learning folks-- if they have some exposure to the topic. It's really, in a way, a matter of creating a common language and culture, and understanding the high-level concepts.

And then, you're able to collaborate much more efficiently, and that's, kind of, what the "Elements of AI"-- and especially, there's a part two called, "Building AI," where you have optional coding exercises, and it goes a bit deeper into the technicalities. And ideally, I would like to see a small company-- maybe there's a CEO who takes the course, and then they have a couple of technical people there. But over the weekend, they've done some parts of the course, hopefully, and they've maybe done some sort of minimal amount of coding.

And then, I can always see when they go on Monday to the office-- hopefully, we'll get to go to the office soon-- they go to the office, and they are beaming and they're bragging, oh, over the weekend, I did a bit of coding. I did this machine learning tidbit here, and it worked really nice. And you know, those bragging rights and the confidence that they get is the sort of thing that I'd like to see that arises from that.

They don't need to be the ones who will write the production code that is actually going to run on the server. But the fact that they've done some coding, and they're not out of their depth when they're discussing it with the coders-- I think that's the best thing. And then, we have that easy understanding between people, and that's going to be more valuable than hiring five more coders, perhaps.

**KIMBERLY  
NEVALA:**

Yeah. I've also just, outside of the corporate environment, been really struck by all of the different paths by which people intersect and become engaged in artificial intelligence. So I believe you also know Renée Cummings, who comes from-- she was a journalist, and then worked in the judiciary, and she's a criminologist and has come- to AI. In fact, she says she invited herself to the AI table.

And I find that so incredibly exciting, to see those folks and their contributions, which are just so important and expand our overall worldview with AI. Both in terms of what we can do with it and perhaps keeping us out of the weeds of what we shouldn't do with it. But have you also seen examples of some of this-- I think what you've in fact called the crazy creative attitude to technology-- and what some of these folks can bring if we just let them in?

**TEEMU ROOS:**

Yeah, definitely. That's been some of the things that we have succeeded in or at least I've felt a lot of joy about seeing some of those examples. One instance was this week when I had a meeting with a person who was taking part in a training program.

It wasn't "Elements of AI." It was another program that we have for local companies. And she was saying that she'd been working on-- now, I don't know the terminology, especially in English-- but creating her own clothing and using the patterns that you can get somewhere and customizing them and tailoring them to fit just right.

And we have this project work that they do as a part of this training program, and she had chosen us as one of her topics-- that application-- so how to adapt patterns so that they fit just right. And she was like, well, I wonder if this is technical enough or something, and I was like, no, no, no, no. Just don't worry about that.

That's exactly what is missing. And this is the kind of example that I would never in a million years come up with, because I don't sew clothes, or I don't even know how that's done. And that's the primary example of me being able to help her identify the technical solutions that might be applicable there.

And she was already thinking, OK, well, what's the business model? This has to be scaled up, and how do we sort of connect people when they're at home? And maybe this is picking up now with people spending more time at home and with people's hobbies and stuff.

And I think that was, kind of, a perfect example of completely complementary spheres of knowledge. I didn't know anything about the application, and she was just learning the basics of AI. And combining them is exactly what I'd like to see more.

**KIMBERLY NEVALA:** Yeah. And that's a good example, I think, again, of harking back to something I've heard you say, that we don't have to ask all of these people in all these fields to become different people to be able to use the tech. We can adapt the tech to make it available and accessible to people where they're at, which is fantastic.

So as you continue this work over the next several years, what do you hope to see, or what can we expect to see in terms of how the public conversation in both access and the importance of education and literacy in AI is going to evolve?

**TEEMU ROOS:** Well, I think there are a few things. I would like to see more active citizen participation in the shaping of even legal environment in which AI is being applied. There's lots of legislation coming out from the EU especially these days.

And I'd really like to see that there are a lot of organizations and individual citizens voicing their opinions-- their, of course, informed opinions-- about those matters. So the public discourse being elevated to a more and more detailed and mature level is one thing that I definitely hope to see. It is hard to tell whether that's been happening, or where it is going.

There is a lot of pull to a certain direction from, obviously, the industry, from the corporations, that they are able to, sort of, lobby for things extremely well. So what I'd like to see on the other side of the table or around the table as well is more the citizens' points of view, and I'd really like that to be more prominent. But there's also the positive, creative side of things.

I'd like to see more innovation, more industrial activity coming from, I guess, small and medium enterprises, so that they have the courage to go and try all things and to dream big about scaling up. And I think that can only happen if we have a really wide basis of people and companies and organizations that feel positive about AI and come up with those really new ideas. So that, I hope to see, and that's something we're really working towards-- leveraging and giving a platform to the people that get enthusiastic and start innovating. And that's really inspiring to see some of those initiatives.

**KIMBERLY NEVALA:** I think that's great. I leave this conversation feeling energized and optimistic and even more aware of the need to continue even my own learning in this field. So if, like me, you are interested in learning more or would like to share that gift of knowledge with your own networks, links to "Elements of AI" and other resources mentioned on today's episode can be found in the show notes. I really enjoyed the conversation, Teemu, Kiitos.

**TEEMU ROOS:** Kiitos. Thank you.

**KIMBERLY NEVALA:** And that's the end of my very limited Finnish vocabulary. In any case, next up, we're going to pay homage to the role of the arts in AI as we talk with Shalini Kantayya about her path to understanding artificial intelligence and investigating algorithmic bias in the documentary film, *Coded Bias*. Make sure you don't miss it by subscribing now to *Pondering AI* in your favorite podcast.