Nemko Digital

Nemko Digital Webinar Report - AI Trust in Electronics Summit 2025 Transcript

---

Bas Overtoom (0:25)

Hello. Yeah, that's nice to meet you. Nice to meet you too.

So welcome everybody. I hear that now we found everybody in the other room and brought them here. So I feel good to start. My name is Bas Overtoom. I'm the Business Development Director for Nemko Digital.

And it's exciting to be here today with you to talk about AI trust and to have a very nice audience of business leaders, compliance leads, academics and government organizations to talk about this topic of AI trust. Because also AI trust is something you cannot do alone. We have to work together to make a step and that's why we are here for today.

Now we focus on AI trust in electronics, but of course many of the things we are speaking about will be relevant for all the other companies. So it's very important to talk about it.

And when I think about AI trust myself, AI is of course something that is not really new, but it is something that is really changing very fast. Already 50 years ago we spoke about AI, but the way it's been changing in the last couple of years is just unprecedented. And recently also China released its vision on AI and they want to have a 100% AI enabled society in 2035, whatever that means, we don't know. But the speed is fast and ambitions of many of the countries are very high on this topic.

So that's why we need to talk about it, because we support the adoption of AI, but we believe that you need to do it in a good way to enable the benefits and to avoid many risks that are also out there.

Nemko Digital

Then about electronics myself, I started at Nemko about a year ago and then we are looking into electronics, AI trust in electronics. And I remember one of the first projects we did was an AI driven air conditioner that could see where the people are in the room, optimize wind. And then at that time I made jokes to the clients to say, yeah, this is just the beginning, we'll soon have AI enabled washing machines and AI enabled ovens, blah blah blah. It was more of an idea.

And this summer, six months later, I was working in Germany in a media market and they were already various products. I could already buy them. So it goes very, very fast and that's why we think it is very important to talk about this topic here with you today and we're also happy with the turnout today.

So that's basically maybe before we go into the program, just maybe a small quiz. So who of you is already actively implementing AI in a product or service that is going out there? So not maybe for your own personal use, say that's what we all do. Then you have 100 hands, but who in the audience is busy with bringing AI to production and to consumers? You can be proud of that. OK, half, half of the people. So that's where we are with the maturity.

Then this is the topic of today and this is where we will focus on. I already explained it a little bit in my introduction from Nemko and from IBM. We believe that the key is to use guardrails to not only keep things safe, but actually to also enable innovation. So if you want to be innovative, it is about building in trust from the start and it can start small when you're just in this POC phase and then slowly grow it over time when you're going into production. That's basically the key theme that we have chosen for today to talk with you about.

And there's a key theme that you will hear again and again in all the presentations, so for the program of today.

Now the first part we already did, that's me and we'll have the introduction from the State Secretary to give a bit of an outlook on the wider scheme and what's happening here, especially in Scandinavia and Norway on this topic. Then we dive into more setting the scene for the industry and giving bit of an overall understanding all the key themes that will be later introduced on this topic by my colleague Pepijn and then we dive into the future on what's happening in New Frontier Agentic AI and then going on regulatory and how it can actually help you to build those guardrails to become faster and better than a short break, and then our

partners from Fisita will talk about actually practically implementing and building AI in a trustful manner, and then we'll try to bring things together in a panel discussion.

And then we hope that you have learned a lot and also we encourage you to actively, especially during the panel, but in between also to have actively contributed to the dialogue.

So that's it. And as a first speaker, I would like to introduce Mariana Williamson to introduce, yeah, her vision on AI in Norway.

Mariana Williamson (State Secretary)

Thank you so much. I don't know guys, can you... Can you guys in the back hear me if I talk without a microphone? Thank you. And also online please let me know because if not, grab a microphone instead.

Good afternoon, everybody. It's a pleasure to be here with you today. The topic of this seminar can hardly be more relevant. AI is rapidly being integrated into systems and devices all around us and changing the world as we know it.

And the technology will obviously also change and shape Norway, just as hydropower, oil and digitalization has changed our society. But to make this shift beneficial for all, we must work together across sectors and industries, and our government's goal is we want Norway to be the best and most digitalized country in the world.

Our society faces major challenges, as you are well aware of: an aging population, labor shortages, public sector renewal, job creation, emission cuts and digital threats, increasingly digital threats. In this context, AI can be a powerful tool for us, capable of solving even some of the problems that seem unsolvable today. But none of this will be possible without trust in the technology.

In a Norwegian study from 2024, only 30% of the respondents had rather high or high trust in AI systems. A global survey by the University of Melbourne and KPMG showed a 58% trust. Though not directly comparable, both indicate significant public skepticism towards AI systems, so creating the conditions for trustworthy AI must be a top priority for all of us.

Nemko Digital

Technology already underpins key sectors in the Norwegian economy, such as energy, maritime, health and finance, areas where Norway holds strong positions and access to high quality data. As these and other sectors adopt AI, we see opportunities for efficiency, innovation and global competitiveness. Norway offers a lot of advantages along the way, and the government has high ambitions for use and development of AI. It is why we are building national infrastructure, working hard to ensure favorable conditions for businesses, and also enabling the public sector to deliver better services.

Our goal is to make Norway a front runner of ethical and safe AI. We must move quickly, but not rush to ensure that Norway's position and our ability to adopt and succeed with AI.

The Labour Party government is taking a lot of measures along different tracks. I wanted to take you through those five tracks today.

Number one is investing in national AI infrastructure. This includes the development of the Norwegian and Sami language models that will be open and free to use for development of AI services, both for the public and private sector. Trained on large relevant Norwegian data set, including also newspaper articles, as we revealed earlier this week, these models will reflect a Norwegian context, including our society, values and attitudes. Including no, sorry, and hence also contributing to the development of trustworthy AI.

Further, we invest in performance computing capacity, a prerequisite for both development and use of AI models. We also know that data centers are crucial for an AI revolution. That's one of the reasons why we launched a new data center strategy just before the summer and we welcome an industry that contributes with jobs, value creation, security and also enables AI adoption.

Number two - We're investing in competence building. We need to make sure that it's easier to recruit ICT specialists both in the public and private sector. That means investing in digital skills from primary school to higher education as well as lifelong learning. Research and development are the foundation of innovation. The government has allocated over 1 billion Norwegian kroners to six AI research centres. One of the centres is the Norwegian Centre of Trustworthy AI, where I know also IBM is a partner.

Nemko Digital

Furthermore, last week we presented the government's roadmap for technology-based business sector, aligning efforts across businesses, research, business research and government. We also have a clear goal of making sure that Norwegian businesses of all sizes can utilize AI to increase productivity and profitability. One example is the Practical Guide for Responsible Use of AI Assistance developed by a group of experts appointed by the government earlier this year. Put it on your must read list because I highly recommend it.

And lastly, we're also establishing AI Norway and I'll get back to that in just a bit.

Number three - increased use of AI in the public sector. As I mentioned before and you are all well aware of, AI has the potential to drive improvements, innovation and efficiency in the public sector and hence also contribute to preserving our welfare states. We've set a goal that 100% of public agencies shall adapt to AI by 2030, and we are putting in place several measures to ensure that we can reach the goal in a safe and secure manner. The number of central government agencies using AI has already doubled from 2023 to 2024.

Or whatever that means, you might say, because there is definitely a huge variety in the way that they use AI, but still a lot of them are at the testing stage and some have now already... already we see several AI solutions being rolled out throughout the sector, creating better services for Norwegians and also increasing efficiencies in the various entities.

In Norway, surveys indicate a strong public trust in the government. To maintain this trust, it will be essential to implement AI in a transparent and responsible manner.

Number four - International cooperation. AI is a global phenomenon and a global value chain, and we must collaborate internationally to address its challenges. Norway is therefore working closely with the rest of the EU, OECD countries and also especially with the Nordic countries to promote safe and ethical AI and also to collaborate on, among other things, competence building and infrastructure.

We're also part of the AI task force for this year's G20 summit, where Norway is once again a guest country.

Number five - We're establishing a national AI regulatory framework, which I know is on your agenda today. By implementing the AI Act in the Norwegian law, we

ensure that Norwegian businesses will have the same competitive conditions as the rest of the EU.

And at the same time, it's also a cornerstone legislation for trust and innovation in AI. As many of you are probably aware of, the AI Act is fundamentally a product safety regulation. It establishes essential safety requirements that AI systems must meet before being placed on the European market, particularly in high risk contexts such as healthcare, education and public services. Like other EU product safety laws, the AI Act focuses on ensuring that AI systems are safe, reliable and respectful of fundamental rights throughout their lifecycle.

This includes requirements for risk management, human oversight, data quality, privacy and transparency.

The draft Norwegian Act is currently being circulated for comment and I encourage you all to provide your response in the public consultation. However, there is no reason for businesses to wait for the AI Act to enter into force before taking action.

Preparing for compliance with the AI Act means your products and solutions will be better, safer and more trustworthy. As a part of the implementation of the new AI Act, we're also strengthening the national guidance and innovation capacity by establishing AI Norway within the Norwegian digitalization agency.

AI Norway will be the government's national arena, a key driver and competence hub meant for both public and private sector. The format and the services they will provide is still in early stages, and I encourage you all to come with input on how to make this arena relevant for you as well.

AI Norway will also host the Regulatory AI Sandbox, where Norwegian companies can experiment, develop and train AI systems within safe and controlled environments. The goal is to strengthen competitiveness and create greater opportunities for Norwegian AI systems. And especially small start-ups and small and medium sized companies will benefit from this.

Guys, trust is essential for businesses. People must believe that AI powered products and services are reliable and safe. Auditing and certification of AI systems can play an important role to achieve this. It's encouraging to see how companies like IBM and Nemko are developing innovative solutions for AI assurance. These

efforts show how collaboration across technology providers and assurance actors can build trust throughout the value chain.

Trust is not only a prerequisite for domestic adoption, it's also a competitive advantage internationally. Products and services built on responsible, transparent and ethical AI will be better positioned in global markets.

Collaboration between government, academia and industry is crucial when setting up the guardrails for trustworthy AI. And as you know, it's a lot of hard work. We all know the saying that trust is built in decades but can be torn down in seconds.

Fortunately, we are all a part of a growing community striving to make AI a force for good. The future will not be shaped by one country alone or by one sector alone, but by all of us together. And to succeed, we must build trust - precisely what this seminar is all about.

Bas Overtoom

So thank you very much, State Secretary Williamson, I think very well said and very inspiring to hear all the initiatives that are already being taken and launched to help here in Norway and further to drive on this topic.

Then going further in the programme, the plan from the land, your hands up please. Pepijn will bring us well, I get your presentation running on the more of a setting of the scene on all the topics and the developments in the industry when it comes to AI trust. I'm going to put the slides right and you can take it away.

Pepijn

Thanks a lot, a lot, Bas, and also State Secretary, because I think, yeah, this is a great moment for me. It's only a couple of months after I joined Nemko. Actually, I come from a background of building AI for over 10 years at a scaler based in Toronto at Deloitte, out of the Netherlands, leading the AI team.

So I'm really from that background of, yeah, let's see if we can make this work and let's see if we can get an impact out of this in the business world. And what I experienced there was basically that all of that impact is hampered not so much by what we can build, but really by the trust that organizations that users have in the solutions. So that was for me the reason to join Nemko and hence also the reason that I'm standing here today.

So let me see if I have connection. And not, maybe I need to be closer. OK, good. So we have connection, we have slides and what is a consultant without a slide then you have nothing.

So basically what we have seen over the last years is a huge increase in the complexity of products, a huge increase in the need of things that we need to trust, right? In the past we had much simpler solutions. So that is really also where, yeah, right now the rubber starts to hit the road in the stage of AI, what does it mean to have trust?

These are 7 principles, 7 principles of the EU. People may recognize them, maybe not. But what you will probably recognize is some of the news items below. There is a lot that goes into AI trust and it's super broad and you also see that not everything is about technology, one thing that's interesting to mention here is...

...intricate systems really trustworthy because you see really the different angles to different dimensions coming together and in the end the whole of this system needs to be trustworthy. So that gives you a frame for the scale of the challenge that we are facing in that.

So it's not surprising in that context that governments are also starting to put in place regulations. And yes, there is indeed the AI Act, which is super important and which defines a lot at the level of AI systems, but there are many other acts and regulations that you need to be aware of if you bring AI powered products onto the market and basically bringing that as a whole. Yeah, that is what you need to be ready for and maybe, yeah, to say something that's maybe from the private sector, not the most expected remark. Yeah, actually a lot what you see in the regulation, yeah, you can say, yeah, it is complicated for businesses to follow it, but yeah.

If you are really honest, a lot of it is also just good development practice. It's just doing things right in such a way that you can actually bring it to a consumer and the

consumer can trust the solution. So that's also the mindset that we hope to bring here to bring those different worlds from academia, EMEA policy together to really think about, yeah, how can we make AI trust work in a practical way and really make sure that it's part of the way that we scale this solution and get the benefits of.

And also good to know is of course the timeline of this. And it might feel a little bit far away because it's 2026 and we're not there yet. But yeah, if you're not started, if you haven't started yet, yeah, then you should really start pretty soon to make sure that by next year that you are ready for everything that's coming, in particular this AI.

So, and this is also where we as Nemko Digital try to find kind of the right balance in helping companies already taking steps towards that compliance and towards that conformity because. As said before, yeah, you cannot wait. So if you just start to think about the regulation at the moment that the rubber hits the road, yeah, you're often too late. So there is also in that plane to there are different ways to look at it and the lens that we have taken there with our AI Trustmark is really saying, look, you have all the ISO standards which talk about your processes, which you can basically get that tick of the box. You have the AI perspective, which for the AI act perspective, which is really focused on the AI system. But there is something very important in between because an AI system almost never lives alone. It always is sitting in a product, in a service. So let's take that as a leading perspective and let's see what do the regulations and the guidance that is already there mean.

For AI embedded products, for AI embedded services and that is what we have, what we basically have started out of asks that we got out of our partners in Korea to really kind of fill that gap between the organizational level and the system level by looking at the product level and Monica will talk a little bit more about that later. In the end, it's all about finding the efficient and effective way to get to that compliance in such a way to enable scale.

Because in the end, that is what we truly believe in. Yeah, you don't do this because the regulator says so. You do this because it has a business advantage. You do this because you're really see a differentiated opportunity for AI.

So that is really, really the starting point. It's also what we see in the in the conversation that we have with business leaders. It's also what I heard in some of the conversations in the at the lunch just before this. Before this event, and a lot of those conversations are not just about how do we get AI to work, but how do we get it to work at scale, right? So that is really, really the core of everything.

So, so in order to get to scale, yeah, we have developed a maturity model where we base, where we have defined 8 essential cornerstones. What you need to do also as an organization in order to make sure that you get your AI in a group state and that you are also ready to scale that AI to full impact and that starts from a leadership in governance level down to also the real the AI life cycle where the development happens and then all the way around to technology and all those dimensions are essential.

And we've developed that into a full framework with 38 capabilities that you need to work on that you need to develop in order to make sure that you're in full control. So that is what we help our clients with and that is also the journey that we see a lot of companies are on at the moment and that journey has several phases and of course in the end that leads to a position where you are also the leader as in AI, but most companies are not there yet.

In particular in Europe, yeah, we see that a lot of companies are still, yeah, most companies have had phase one where you do the exploring, right, and are moving into phase two where you have scattered AI initiatives, where you have people around the organisation going off and do their thing often enthusiasts, and this is one of the most risky situations actually from a control and from a scaling perspective, because everything that you basically have built in that phase and without the proper governance is something that basically can slow your development later on in the process. So this is where the shadow AI starts to emerge and this is where you really as an organization need to think about how are we going to do the governance, how are we going to build that trust and the earlier the better because the less you have to pay.

And that's also, yeah, leads to a different way of thinking about AI maturity. A couple of years back, I was talking to people in the AI space about AI maturity. They would also always talk about the AI development lifecycle. And they will talk about making that better for developers. But that's not enough anymore. We have seen that AI is coming, is coming to the market. So that means that you need more

than just thinking from a developer perspective. You also need to think from a governance perspective.

And really make sure that those operational level of those organizational level controls are also in place. That's also why I'm happy that we are doing this today with our partners from IBM who are heavily investing also in their technology solutions to not only enable this kind of thinking, but also enable it at scale. But we'll probably hear more about that in the following presentations as well.

Maybe then to round off a little bit more, look into the future to what do we see coming at the moment. Agentic AI, of course, very important development. I'm going to make it relatively simple for you. Just look at the side of the screen.

There you see basically the four key elements what makes AI agentic. With generative AI, we have already seen that it's bloody good in interaction, in communicating that you can connect it to a knowledge base and then you have a nice chatbot solution.

What we see with the new solutions, the Agentic solutions, is that basically there are two capabilities that you add into the mix. One is that this AI is not only sitting there on a bunch of knowledge, but it's also empowered to take actions. It's connected to other systems. It can write to your database. It can really change things in the real world. And the other thing is that it's not just AI doing things on behalf of human now, it's also orchestration. So agents are basically steering agents, so they are.

You also get a lot of new risks from those two capabilities that you add. We've mapped that a little bit on the other hand of the slide where we already saw that there are new AI risks emerging with generative AI, adding those new capabilities. Also opens a new volume of risks, maybe to mention two. If you have AI agents talking to AI agents, yeah, there's a serious risk of errors that are in the system being amplified before they get caught.

And the other risk to highlight is the harmful actions. Yeah If this AI starts to do actions, starts to write to the database, yeah what does that decision that the AI take mean for humans and where is it harmful? So we see that this is still a very fast evolving field. We have now, I think, a number of frameworks which are here to stay, which allow organizations and companies to really take steps in AI trust.

But yeah, it's definitely only the beginning of, yeah, in my perspective, exciting journey. So next step in that journey will be the next talk by HP. So before we, yeah, before we do that, I'll hand the mic back to Bas.

Maybe to summarize the talk in one sentence, yeah, stop bolting on trust as a feature of AI, but really build it from the foundation. Thank you.

Bas Overtoom

Thank you. The time. Maybe one minute you can stay here. You're not that easy. So yeah, overwhelmingly, a lot of things are going to happen. I can imagine you sit in the room, you think, OK, my God, we're just getting started with AI. There's so many additional things.

So going back to the, let's say standing in the shoes of let's say just a business leader that is want to do something good with AI, what is the first or maybe the first two steps just to get started with? What will be your advice there?

Pepijn

Yeah, I think if you're, yeah, I think in most companies, if you are a business leader and you haven't gotten started yet, there will be a couple of people in your organization who have already started and who are basically tinkering maybe even in the evening hours to see how they can make their job easier.

So my first tip would be really find those people, talk to them and make them your ally, because that is really the opportunity where you create a lot of speed and also can teach those people that it's not just about building nice stuff, but also building it in a responsible way. And that way you can really kind of yeah build a partnership between the trust side of AI and the development side of it.

Nemko Digital

Bas Overtoom

Thanks. Thank you. Is there somebody in the audience have a particular question to prepare? You can keep it also for the panel. Thank you very much for that. Thank you. All right.

I think we spoke already a little bit about Agentic AI little, little, yeah, a little bit of an introduction. But now we will sink deeper into that very, very important topic for the coming 20 minutes. What is happening with Agentic AI? And we had generative AI. Wasn't it already enough? No, it is just starting point is only the beginning of more to come HP.

[Applause]

I will get your start ready. I'll get you started.

Where's my beautiful selfie? Sorry, it's my beautiful. I know there's my beautiful. Thank you for spending an afternoon with us. It's really great to see so many people wanting to spend an afternoon with us here. I guess my claim to fame is that Inga Strom cast me to write a forward in her book. And thank you. And I did that. So yeah, got me. I lead IBM's AI for business in Europe, Middle East and Africa, original in Norwegian, but I've lived outside of Norway for many, many years.

HP (IBM)

How many here are you seeing Microsoft copilot? See, that's not a surprise. There's a reason why I asked. Take off.

That should work. Thank you.

Personal productivity. Does that make you AI powered? How can you calculate return of investment on personal productivity? How does personal productivity tools give you a competitive advantage when everybody else is using?

At the moment, the maturity above personal productivity is fairly low, actually from an IBM perspective, and I usually say that IBM is either one of two things. We're the world's oldest start-up. Or the world's oldest AI company. Both are probably true.

We've operationalized about 6500 AI systems internally in IBM. We operationalize around 200 new ones every quarter and our mission is not personal productivity. We use Copilot as well, but that benefits me to hours of my time say is 2 hours of my time, so doesn't help IBM unless I contribute those two hours back to the company, right? That's what personal productivity is really, really difficult to calculate. I do spend a little bit of time talking about use cases because that's why AI is so fascinating.

As a technology, it's been around since sort of late 50s. We've been involved throughout sort of every cycle of AI. We've seen decades of heavy investment, decades with no investment. The 80s was a barren decade for AI in the 90s too, because of the Internet.

But in the 90s, actually IBM developed a computer that beat Gary Kasparov in chess and being the first machine that did so. And the current wave of AI was fuelled by IBM's Watson technology that beat two reigning champions in Jeopardy in 2011.

And that sort of kick-started all of this, oh, natural language generation and machine learning. That's a very powerful combination. And since then we've been involved in about 50 thousand, 60,000 AI related projects globally.

So AI didn't start in 2023 when OpenAI launched generative AI and the ability to talk, you know. But right now everybody says it's AI powered. We see AI washing everywhere that people, companies are overstating how much AI they use.

When you hear, well, 94% of companies in the insurance industry use generative AI on a daily basis, yes, but it's personal productivity. Very little is around core business processes and actually very little has changed in the last year and a half.

A recent MIT study from last week suggests that still 95% of generative AI sandbox experiments fail to get operational.

My take on this, and this is based on a lot of conversations with business leaders, is we don't focus enough on the use case. We don't focus enough on the return of investment of the use case. One example is in Denmark where I live. I'm not going to mention name. There's a big supplier of technology to the public sector.

Thank you.

They suggested to the Danish Railways that they spend a couple of millions trying to find out if there was any customer satisfaction in when a train is 15 minutes delayed in Roskilde station and it may have, you know, gotten that time back by the time it hits Copenhagen to use AI in some mysterious fashion to inform the clients up front and everything more. We're spending millions exploring that. That's not return of investment case. Where's the return of investment?

That's just experimenting for the sake of experimentation. So if we want to succeed in this space, and I know I've been spending an awful lot of time on the first line. We have to move. There's a road to go. This is IBM. After all, we have to start looking at where is the return of investment and start experimenting on use cases that makes business sense for core business processes, right. So where are we?

Up until 2022 AI can predict it for you. We still do that extremely well by the way. And just before the summer I was in a AI scale for public sector sort of afternoon thing in Copenhagen where one of my favorite use cases. Their scale is an insurance company and actually using machine learning algorithms to predict who has a higher likelihood of getting long-term sickness, things like depression or stress related to symptoms, etcetera. Because that's very expensive. It's expensive for the person. It happens too, but also the companies and the societies and everything. So what they do is based on a certain number of historical facts, very small machine learning model. They've identified so far 3000 people at high risk and then they use a health professional to do what they call a proactive care call.

And 97% of those 3000 were very, very happy that they received a proactive care call. And it has reduced the likelihood of people in that group developing long-term sickness by something like 76%. It's absurd. They thought maybe it would between 8 and 10%, but it's 76, that's such a high number and it's...

...are laid down in the law. If there are new developments, you always have to open up the law again, getting to all of these legal discussions, the entire process, Parliament, European Council discussed, the European Commission making, all of this is not necessary.

Because here you have updates, regular maintenance of the standards, so you can include new technology into new versions of the standards and thus also respond to innovation in the market. So people sometimes ask this question is this AI act not outdated by the time it gets in and get it by the time it applies, we should now say.

And I said no, no, it sets some very high level legal objectives, but the details they can be determined in the standard. It does not go into technical.

OK, and this is laid down just so that you believe me in Article 40 of the EU AI Act, where it clearly says high risk AI systems or general purpose AI models which are in conformity with harmonised standards shall be presumed to be in conformity with the requirements set by the respective sections. So this is the key article to put the AI Act into this context of the EU new legislative framework.

Dating back to the days of Shaktivel. OK, now we have these three areas that are ruled out by the AI Act. You have prohibited practices. This is already in force since February the 2nd of this year. You have the rules for general purpose AI models. This has been enforced since the 2nd of August and I will talk about this, what this means in a second, how you are compliant with this and you have AI systems in high risk areas and here we have a bit less than a year still to go.

Until this applies and there are some other regulated areas, we don't need to look at them later down in annex one, they have one more year to go, but these are the areas that are already heavily regulated financial sector for instance, they have a bit longer here to adapt their regulations to this.

OK, so this is where we are. I introduced a bit the difference between the legal process and the technical process. We are somewhere here. General purpose AI already applies iris systems. The standards are in progress. You see it here at the bottom.

Does that work yet? Very good. OK, so for general purpose AI models, European Commission did something they wanted, you know, when they finalised the regulation, the AI Act, sort of six months before in the middle of the negotiations before finalisation.

The ChatGPT appeared and everybody was like crazy using ChatGPT. And of course the politicians all said we need to regulate this as well. We need to do something. Yeah, this is not in our scope yet. And then they debated, debated, debated. And then they came up with these rules for general purpose AI models. But they also said this is so urgent, we need to do this.

And what they decided is here, they said we cannot have the harmonised standards in time. So they said we do something like a code of practice to bridge

the time until harmonised standards are available. So this is where we are today. Whoops. So this is where we are today here in terms of the two processes.

Now a lot is currently being discussed. If you follow this, will the AI act be delayed? Yeah, there have been 20 companies, their CEO sending an open letter to the European Commission that it should be delayed. It's the and this applies for the high risk AI models, high risk AI systems.

OK, all indications we have so far are that it will not be delayed. I saw that Nemko even have it on their website. You can click here. I provided the link here. You follow this very closely and you also say there is no indication on LinkedIn.

There was from the spokesman of the European Commission, Thomas Renier, a clear statement saying we take the concerns raised serious. Yeah, of course they take the concerns very serious. What else should they say? But illegal tax.

Is a legal text, and in order to change the deadlines or whatever, the legal text would have to be opened. Yeah, the timelines are laid down in the law. Yeah. So what the European Commission would have to do is make a proposal.

To change the timeline that were painfully agreed 1 1/2 years ago and in long discussions. OK, but the discussion continues. Nobody knows where it will go. So just yesterday there was leaking information that the Czech Republic.

The Commission saying they should open the timelines, postpone the timelines. And today there was also information that the Member States in Europe who need to set up the market surveillance authorities will probably not make it in time.

So this will also be a problem. Discussion is to be debated, but bottom line, for the time being we can assume two things.

The legal requirements will not change. They will be the same. And whether they come into place August next year or in two years from now, let's better prepare for them. Let's not wait, let's prepare to fulfill them. And secondly, if some delay is decided. It will not be decided until early next year. It takes time to discuss that, to prepare it. OK, so general purpose AI models. I already mentioned that the European Commission said we have this code of practice to bridge. This has been developed. It's available. IBM has been an early signatory to it. So there you can see on the Commission's website a list of the companies who sign it.

We fulfill the requirements in this. Like many other companies, it's about two dozen of companies, something like that, who are ready signatories. They require openness information to be published, and as I mentioned on the long run, this will be replaced also by a harmonised European standard. Here just a little glimpse the two key parts of this code of practice.

There is a part on transparency and a part on safety and security. The part on transparency is very straightforward, just has these three measures drawing up, keeping up to date, model documentation, providing relevant information, ensuring quality, integrity and security of information.

That's fine. The security and safety part is a bit more complex as you may expect with a topic like security. It has in total 9 commitments within specific sub measures. Yeah, I just copied so that you get an idea what is being addressed there. The first one, the safety security framework. You create the framework. Implement the framework, update the framework, framework modifications. OK, so a bit more complex, but still it's doable and actually every good AI technology should follow this. Yeah, OK, so IBM does have a couple of models out our IBM granite.

Series, for instance, that are all available on open source. And yeah, so just to give you here the advantages to a bit of marketing the open source, you can choose the right model from sub billion to 34 million parameters. It's available.

Under Apache 2.0, it's performant, it's trusted. We do have all the openness information documented. We are in the process of putting them also into the template which the European Commission published and which they recommend, but more or less it's the same kind of information that's available and published.

So it's doable to comply with this. OK, whoops, wrong button. Now to the second part, high risk AI systems. This is a bit more complex. Essential requirements will apply roughly in a year, a bit less than a year from now, 11 months.

What is currently in process is the development of the harmonised European standards, which I mentioned according to the new legislative framework standardisation requests. As I said, once they are available and correctly implemented, they will provide the presumption of conformity and no third party.

Certification is required, but a so-called conformity assessment process. These standards are being developed in consensus with all stakeholders participating in

these two of the three European standardization organizations. So you have three formally recognized European standardizations.

Standardisation organisations, CEN, CENELEC and ETSI and CEN and CENELEC have a joint technical committee, JTC 21, and this is where this standardisation work takes place. Usually something like 200 people coming together discussing, you can imagine.

Finding consensus amongst 200 people who don't just represent themselves, but represent organizations, and it's academia, civil society, research and industry. It's not easy. It's not easy. I could. I could now spend the rest of the day to complain about the process.

Complain about the delays, the discussions here, but we need to work with this. We are used to working with this. IBM is actively involved. We make proposals, we try to drive this, we believe into this system because it's good. It's lean as HP said. We have it for the harmonized European market if you compare it to the US where you have a complete.

Fragmented, completely fragmented legal framework. This is, this is beneficial. Yeah, it's also good for industry. You rely on the experts to contribute. Yeah, but discussions need time. Discussions need time. Finding consensus needs time. I will come to the issues later on in the slide, OK.

There is also an international level. This is maybe also very interesting for a country like Norway in ISO and IEC, two of the international, two of the three standardisation organisations recognised formally international under the WTO.

There is a Joint Technical Committee 1, Subcommittee 42, which does AI standards. And of course, for globally active organizations, there's always the ambition to have the international and the regional specific, the European standards.

As closely aligned as possible. What we currently can see is the European stuff is at the forefront. They need to be available because the legal timelines are the legal clock is ticking. But on the medium long run we need to find a way. Also Europe in general as an export oriented continent we cannot.

We cannot afford having specific too many specific rules in Europe. We want to have these standards identical, which then also, by the way, is the Brussels effect, which in Norway know especially well, but which also the US, China, Brazil know if companies need to fulfill the strong European requirements.

They will do so globally, yeah. And then they export this into the rest of the world, yeah, without any other regulation needed. OK. So what is also doing a conformity assessment? I spoke about it member states. Member states are currently in the process of.

Building up their market surveillance and I just mentioned they seem to send alerts that they will not be able to do this in time. It's a funny thing. Yeah, they they also have known about having to do this for quite a while and they are doing accreditation of the notified bodies with which you can do conformity assessment.

Notified bodies like Nemko is one for several areas of the hardware of the regulated work already. OK, so looking a bit closer at the legal requirements, what is doing on in standardization you have on the left hand side.

The essential requirements in the law, they say you need to have a risk management system. You need to have rules for data and data governance. You need to have technical documentation, record keeping, transparency and provision of information to users.

Human oversight and you need to have rules for accuracy, robustness and cyber security and they get translated into a request for 10 areas of standardisation. And it's no surprise it's more or less a one-on-one equation. So the standards being developed are on governance and quality of data sets, record keeping, transparency, information to use.

Process human oversight on accuracy, on robustness, on cyber security, then a quality management system which is required according to Article 17 of the AI Act and the conformity assessment process.

What is not yet done in standardization is the work on general purpose AI models. The Commission said. Let's work with this code of practice, as I mentioned, let's not interfere with the current process because yeah, there's only limited number of resources, a limited number of time and sustainable AI. There is a requirement in the AI Act also that AI systems look at sustainable AI. This is, this is to come.

This is the high-level architecture. It doesn't give much additional information, but how these standards are being developed in terms of an architecture. So you have a so-called trustworthiness framework that covers some of them and you have the standardisation where your battery is running low.

Oh, Oh my goodness, you solved that. That's good. OK, yeah, so you have these different standards that are being developed. Not much use, just put into some graphic. And as the battery is running low, I probably need to push the button here in order to move on. I hope I don't screw anything up now.

No, I do. Oops, I need you because this is not maybe it's only this battery which is running. OK, is this OK, good. So to look a bit deeper into this and I told you it's complex. Yeah, this standards development and you don't need to.

Try to read all of this. It's just to illustrate a bit of complexity. So you have risk management, quality management, you have data quality. Then you have here a group under the trustworthiness framework, record keeping, transparency, human oversight, another part of the trustworthiness framework for accuracy and robustness, cybersecurity.

30 years.

Special one and conformity assessment. And yeah, you also see that some international standards which are available are being considered already to be referenced, to be used in order to avoid duplication of efforts. All of this is going on. It may help to understand a bit when I say there is a lot of complexity and there is.

OK, these are the current timelines and to be honest, I'm not even sure whether they still stand, but what you also see is 9th of July 26th, 9th of August. Yeah, August 26.

It's getting very close to the final deadline when all the safety routes apply and what is already clear now, some of these standards will only just be available in time. Not yet be available. They will definitely not yet be listed in the official journal, so this will cause some.

Issue, yeah, because you will have to be compliant. The standards are not there. And if you get a standard only in August, I mean, we at IBM, we know what's going on. We know the standard, we participate. But those companies who don't

participate, they don't know what's going on. They will have to wait until they can purchase the standard from the respective national.

Body who are the ones who sell them, you need time to implement it and you cannot implement it in two days or whatever. So there will be some issues that need to be solved. I think we can skip this.

This is just one idea how the cyber security standard is being addressed or maybe the controls are interesting. Controls are that are being looked at main requirement rational.

Applicability, exceptions to applicability, sub requirements, recommendations, risk reduction, guidance. OK, just a little postcard here of what is going on. The situation today is starting point.

Harmonized standards are under development. At the same time, everybody needs to consider it's a new territory. It's not the traditional product safety standards. Going into AI was a new territory. It took an awful long time for all the experts, for everybody to say what do we need to do here to make these standards?

The situation currently is the process needs time. Some standards may be late. Many drafts are currently too complex. Still they are. Not implementable by industry. So there was sort of collection of input of input of input to the standards. Now the task really needs to be throw everything out again. That is not necessary, yeah, to have a standard that we all in industry can implement.

That you can test again and avoid over complexity and then the risk. As I said, standards may not be available in time or not listed. And here market surveillance in the Member States and the various countries need to give a clear statement, clear direction for this highly likely situation.

What we are waiting for something this autumn to discuss with the market surveillance authorities. And this is not easy either because...

...in the healthcare.

So just to set the scene and give some background information. So currently doctors spend a lot of time going through medical documents searching for relevant information to base their decision making upon and this search for relevant

information is done manually. It's slow, often depends on the doctor's expertise to know where to look. So this increases the risk of inconsistent decision based on the doctors expertise.

So what we did was that we sat down with the doctors and had a workshop and then mapped their workflow and identified areas where we could improve the process. So the outcome of the workshop was that we decided to test a RAG system that could support their work in finding relevant information and thus reducing their time spending on this.

So I just want to stress that this is a support tool. It helps find relevant information, doesn't replace the decision making, but that responsibility lies with the doctors.

So here we have a simplified version of the workflow, so many lines of work. You need a medical certificate in order to do the job. But if the doctor deem you unfit to do the job, you may appeal the decision, and if you do so, the appeal goes to a committee and that committee will then decide whether to overrule the decision by the doctor or if the decision stays.

And in order to do so, the committee has to go through a lot of documentation in order to find relevant information to base their decision on. And this could be all be, including clinical summaries, work descriptions, prior decisions and different guidelines. A lot of documentation to go through.

Once they have gone through the documentation and find all the relevant information, they will review it and then make their final decision.

So as I mentioned in the beginning, this process of finding relevant information is time consuming. It relies on the doctor's expertise, so there's a risk of inconsistency. And risk missing essential information. So what we want to help them with is to save time so they can spend more time with the patients and give a baseline to have more fair and consistent decisions.

So the solution we ended up with is a RAG, a method of retrieving information from a database. So based on the case that the doctors is working on, this AI system will then go through all documentation and retrieve the cases that are similar to the one the doctor is working on and then he will present this information in a nice overview so the doctor easily can extract the information that is relevant to make his decision.

So it will. So the doctors will get relevant information faster.

Here you have the proposed workflow replacing the manual search with a AI search.

So when developing this proof of concept, there are has been some challenges as we are working with medical data. There has been some challenging issues with getting access to the data we need. Data in order to build the system and verify that it works as it should. So getting access to data and processing it, removing personal information so we can work with it has been time consuming.

But necessary, and similarly as a medical data, we need to ensure that the whole system is on a secure platform. And then we have the quality part. We need to ensure that the system has the is working as it should, that it doesn't miss any essential information and that in ensuring quality we will build trust. Are we the doctors or the doctors trust this system?

And the doctors need to be able to verify sources. Where are the information coming from? And then finally we have their business. We need to prove that the this actually helps the doctors.

So this is maybe the most technical slide. So here you have the architecture and data flow of the solution. So it's consists of two parts, one feeding the document and storing the documents in a database. And the second part is how to retrieve the relevant documents. So in the first part, the original data are passed through the identification pipeline.

Once that is done, then we divide each document into four different topics, which is decided together with the doctors. So each chunk or topical chunk here represents a piece of information. That the doctor want to answer. So for each case they have, there are four questions they need to answer. So we divide the documents into topics based on the questions and then we embed them making into vectors so we can store in our vector database.

Once that is done, we can start retrieving information that we want. So together with the doctors we have 4 predefined queries based on the information that they want to have answered. Then we embed these queries into vectors. Then we can do topical filtering and then do a vector search and then we can retrieve the

relevant chunks for each query and then present the information in a dashboard for the doctors.

So we are on a POC level right now, so we want to scale this up and then we want to do this. There are important steps to we might need to do.

So the first one is the data and privacy part. We need to scale up the identification pipeline to production. Then we need to ensure that the performance we have on the test case that we are doing now is still the same when we scale it up.

Then we have the governance and compliance part. So we need to monitor the system and then we need to ensure that we are compliant with healthcare rules and EU act.

And finally, we have the integration adaption part. We need to ensure that the system actually fits the daily routine of a doctor so that that everyone starts using it.

Yeah, even though it's scary with all the regulation, you can't get anywhere if you don't start. So starting, adapting AI, starting with the use cases that give us value and getting ahead. By getting more familiar to where we should go.

Panel Discussion

Bas Overtoom: OK. Thank you very much. Two very relevant cases then. Yeah, I think I would like to ask you to join the panel, Erik, Runer, Pepijn and HP and then we go into let's say more than the interactive. But also here on these kind of things. So yeah, we have some familiar faces but also a new face. So Erik, I think I would like to you're from innovation Norway like to.

Then I'm confused again. But then you can explain me exactly where you're from. If you're not from innovation, Norway, then. So the key question is, who are you? What are you doing here? I went to the wrong meeting. No, I'm.

Erik: And I'm from Digital Norway. So Digital Norway is a company that is set up by sort of large Norwegian enterprises. And what we do is we make courses, we make educational material for anybody that wants to adopt any kind of digital technologies, which means most of Norwegian employees. So our main target

groups is established SMEs. So not the likes of those that develop systems, but those that use systems or hopefully use the system because they don't.

So our main goal is to increase the uptake of digital technologies amongst Norwegian enterprises, companies, mainly small and medium size and gradually also on side of let's say public sector companies.

So yeah, myself, I'm not sure what I'm doing here because I have no knowledge about electronics. I'm a naval architect. Electronics is very it's too small for me. I like big ships. But I have extensive experience from working with companies and I have extensive experience in actually deploying technology. So I think that's where I can maybe have some perspectives on how this can be useful for companies, thank you.

Bas Overtoom: And then I think the name, it was quite close. Sorry, not so bad that I was a bit confused. Great to have you on the panel and since you are now joining as the first one, I think the first question also for you, you've been sitting in the audience, you've been listening to a lot of talks, many different perspectives. Is there something that wasn't mentioned that you think, hey, this is something that I would have expected to hear a bit more about? Or you can also choose for the other option that is, what is the most important thing that you hear? That's either I or B. So wasn't missing something or what do you find especially relevant?

Erik: Thank you so much. Maybe one thing I missed a little bit in the discussion of the AI Act is it's actually quite different whether you are a provider of a system or you are a deployer of a system. I think maybe a little bit more about that because most of you are probably providers of AI system, but if you are a deployer, in other words you use the system. Actually the AI act isn't that scary and secondly.

If you use it for high risk applications, there are ways to sort of, at least in the beginning, use it for sort of what we call transparency requirement, which is more voluntarily making it easier for company to actually start AI system without having to bother too much about the AI system. So a little bit more navigation help in that sense would be helpful.

I really liked and I was confused. Do you work for HP or are you called HP? Yes, exactly so. And it took me and it took me some time to realise, but I think your point

in the difference between the different applications of AI, for instance generative AI versus what I call analytical AI in my view.

I think we have confused Norwegian enterprises by talking too much about generative AI. I think the main benefit for most Norwegian companies given our way of doing business is actually analytical AI and we must not forget about that. Thank you.

Bas Overtoom: Thank you for the reflection. Is there from the audience a question or a reflection for this panel already? Yep, there's the first.

Sanjay Mishra: Yeah. Thank you, Erik. In fact, to know that you are responsible for the improving the digital digitalization in the Norway, right. I'm Sanjay Mishra. I have a question that it's OK that you are trying to build a very strong digital system for majority of the companies here in Norway.

But what do you think that how are you going to secure that digital asset of this country? Secure means when you are depending a lot on the foreign, on the technologies, these are based on the foreign technology. I can't see here that what the digital lobby or the government, they are working to secure their digital assets. My question is related because I'm also a general chair of second International Conference on Digital Sovereignty and in digital sovereignty we are talking about ownership, privacy, security and trust.

So I have not seen. In fact, I'm also working as a researcher on these four aspects of the digital assets of the Norway. I'm working here from the since last four years, so but I have not seen the politicians. Companies, they are talking much more that OK, we have to secure, we have to secure our data to inside this, but I can't see anything where this solid framework or implementation on this stage. Thank you.

Bas Overtoom: OK, thank you. Now very important. I mean we can do a lot here in Norway, but we are dependent also on a lot of other, let's say, fundamental technologies from all around the world. What do we need to do about that? Anybody a volunteer in the panel that wants to pick that one up?

HP (IBM): I see one. It's nothing. That is always a good one. I start. Did I nod? I thought I was trying to hide. I was giving it to you, of course. Well, I mean, Nemko are the experts. No, I think Jochen's presentation tells us the importance of

involving industry as well as regulations and working across countries and actually continents as well.

We call this today for trust in AI, right? I had a very interesting conversation with Gartner the start of the year in this world. Trust is only important as long as you're an immature organisation when it comes to AI once you become mature and you fix your trust issues through governance framework. Risk frameworks or whatever, you start worrying about your target security. How secure is your data that feeds your AI system? How? How can you avoid things like shadow AI, shadow agents and prompt injections and all that kind of stuff, right? And.

I don't think we're necessarily at a maturity level yet where we can talk about that, but we should, you know, absolutely because there's so much vulnerability that it isn't resolved with governance necessary. But we need completely different systems. We need AI and cybersecurity systems to coexist with AI governance systems and that's sort of the direction we see. But this is only possible through you know international cooperation. We need to have some common security standards as well which you know tell that.

Pepijn: You want to elaborate on it to try you, but you said put in it. I think it still works. I think excellent, excellent point. I think maybe. Yeah, 2 reflections to add for me. The first one is already we are seeing this geographical battle being played out in particular, I think around intellectual property rights when we see which companies have signed the principles for and which not? I think the major differentiator there is which companies are willing to take also responsibility for the data that got into the systems, I think that's a very important one.

The other one that's that I think is very interesting is that am I affordable? Yes. OK, I will get that. Yeah, something that also always is interesting is that everyone says that European legislation is the most strict one in the world. I think also with what we see happening in the US when it comes to World AI as they call it. We can no longer say that. It's just different priorities, the different governments and different societies will put only on and that will also be something that you will have to navigate.

I think where Europe excels is introducing bureaucracy.

Nemko Digital

Erik: All right. And I made them. I made the mic work again. You want to add something? Yeah, just just a thought, maybe related or not related. Not sure actually, but but if you look at how Europe is making regulation, how US is making regulation, they are certainly influencing each other's and with a global industry like we have now, I think, and also the trade situation going on, I think decisions made in Europe may be influenced by other things than the regulations.

So for instance, Europe can be forced to actually approve or adopt the sort of business practices from American companies to prevent customs. So we're certainly in a situation where your question is very timely, it's but it's extremely complex and Norway has very, very little role in this because we are so far away from Europe, we are apparently even further from the US. So or say is is really that that we have to adapt. Are we doing it well? Not so sure, really not so sure. And we should really have the politician, the state secretary would would actually be the right person to address your questions. I can understand her. She knew these questions were coming. I asked her you want to join the panel, but not a meeting.

Bas Overtoom: Somebody from the audience from another reflection or a question. Oh, there's two. Now I go on the 1st row here.

Lev Vonensen: So thank you very much. My name is Lev Vonensen. I have a question regarding building a system. I mean you're working with a very complicated system. We saw the example from the something from Bergen here, and I guess you are not building it from the ground up. You're basing your technology and your system analysis and all that on other systems, so you have to rely on other building blocks so to say and so my question is how can you? You're sure that the building blocks that you're using and even more if you use the for other regions like China, US and so on, how can you put the assurance that you're compliant with the with the regulation. I think it's kind of hard. Yeah. So burn building, you're building up on data on other systems. But yeah, governing you cannot go how deep can you go into the system. So it's maybe a bit on the build side, on the governor side you were going to give it a first go.

You know how you do it and maybe you can share a bit later on. I don't think it's that different if it's an AI system or if it's a digital data flow with other systems. You have the same issues. You have the same...

Nemko Digital

[*The transcript continues with more panel discussion and Q&A, including questions about IBM vs OpenAI differences, trust in AI systems, security concerns, responsibility and accountability, business value vs productivity improvements, and concludes with closing remarks*]

Closing Remarks

Bas Overtoom: And thank you for your beautiful questions. I think we can continue for an hour asking more questions. And the good news is we can, but then not in this group setting, but in more individual settings and then we'll be back to the other room.

And it will be some refreshments there and it will be a moment for you to, yeah, reflect with people you know or people you don't know on some of the things that have been discussed. And but you're also at one point excuse to leave if you have your other things. Thank you again and wish you a good evening.

Thank you.

End of Transcript