



CLOUD REALITIES

CR022

The problem with AI with
Theo Priestley, Futurist

CLOUD REALITIES



LISTEN NOW

Capgemini's Cloud Realities podcast explores the exciting realities of today and tomorrow that can be unleashed by cloud.

CR022

The problem with AI with Theo Priestley, Futurist

Disclaimer: Please be aware that this transcript from the Cloud Realities podcast has been automatically generated, so errors may occur.



[00:00:00] I think we're done and, and just I'd like to point out that I did that in a Warner. Everybody. Yeah. Alright, Dave. We'll forget about that though. Sorry, I wasn't paying attention. There was a squirrel outside. Oh, sorry. That, that, that threw.

Welcome to Cloud Realities, a conversation show exploring the practical and exciting alternate realities that can be unleashed through cloud driven transformation. I'm David Chapman. I'm Sjoukje Zaal, and I'm Rob Kernahan

In this week we're gonna be talking again about the subject of the moment, but this time looking at the concerns around AI and some of the ethical dilemmas.

It's growing in volume, the noise around this following the release of chat G P T four, and an open letter signed by Elon Musk and hundreds of others urging the world. Leading labs to [00:01:00] pause training of the new super powerful systems for six months saying that recent advances in AI presents, quote unquote a profound risk to society and humanity.

So it's a big subject and, and joining us to try and get under the skin of. What's going on there and how big the concern really is is Theo Priestly a futurist advisor known for his forthright views on the negative impacts of technology on society, and he currently works to inform political business leaders of the risks and benefits of adoption.

So Theo, really good to see you, Jonna. Just introduce yourself and say a little bit more about what you do. Yeah. Thanks Dave. Hi everyone. My name's Theo Priestley. As Dave said, I am a futurist. I've been in the technology industry for a little over 20 years now in various guises working across business and enterprise software and consumer software as well.

I do a lot of trend analysis and reporting. I spend a lot of time writing about these trends and also trying to examine what the. Positive and negative impacts [00:02:00] on business, technology, and society. These emerging trends have, as well as doing some keynote talks at international conferences. So let's start with Musks letter then, and there's a, there's a really choice quote in the middle of this that I thought was, was worth just reading out to give us a sense and a, and a grounding in, in what's going on with AI at the moment.

So, quote unquote, advanced AI could represent a profound change in the history of life on earth, and should be planned for and managed with commensurate care and resources. Unfortunately, this level of planning and management is not happening. Even though recent months have seen AI labs locked in and out of control race to develop and deploy ever more powerful digital minds that no one, not even their creators can understand, predict, or reliably control.

I think there's a kickoff. Theo, is that a fair representation of the current state, do you think? I think it's it's, it's an interesting. Time and it's an interesting time to release this kind of letter, especially now that [00:03:00] the, the, you know, the main people or the main supporters of that letter el, including Ellen Musk, have decided to launch their own.

AI company after calling for everybody else to stop work for six months. What? Yeah. I think what we need to sort of understand is what do we mean by advanced ai? Where we are is we have artificial intelligence. We have very clever algorithms and machine learning that's taking place here in large language models are just an extension of that.

And we're moving towards a world, I would say, in the next sort of five to 10 years, where we have a G I, which is Artificial General Intelligence, which is an AI that can, that is almost on par with. Human intellect and and human capability to be able to fully realize that you have to marry it with a a physical body, which is why I think open AI is invested quite heavily into one X, which is a robotics company, and you can see the path that they're [00:04:00] trying to



chart, which is embedding a G P T model or a large language model into a humanoid figure. So you can actually talk to it and converse with it and do things, and it can take instructions. But where we are just now, we're still, it's just very, it's just a very clever algorithm. Mm. But the level I think of automation should, should be something that we're concerned with more than it running amok and, and destroying the world by stealing nuclear codes and sending off bombs, right?

Uh, which is everybody's kinda sort of favorite scenario that they like to play out. I think the more concerning aspect here is the level of disruption to jobs, which should be taking front and center here, and also some of the unintended, unethical uses of this technology. So, I've seen people using this as a companion app, as a relationship app, you know, something to have almost sexual relations with, or sexual conversations with your own personal therapist. None of this is something that I. A chat [00:05:00] bot should be used for. Yeah. 'cause it's not an emotional thing by any means. You do not have a relationship with something. It does not convey empathy in any way because all it's doing is, and it doesn't understand the context of what you're talking about. It is just literally a language model that strings sentences together in a very clever way.

Mm-hmm. But go back to the first point in terms of the job disruption. That's the thing that we should be really concerned about at this point. So when people talk about risk, to your point, it's not the singularity moment where, uh, it all of a sudden becomes self-aware and then start to take over the world.

It sounds like we are, we are somewhere off a moment like that. It did seem to me that some of the. Hyper ball around it, even even connected to that note was suggesting that that might be a higher risk closer to home than I was expecting it to. And the bigger risk, and I think this risk is actually now.

You know, actually manifesting itself, it's a real issue, isn't it? [00:06:00] Which is around the job market and looking at how disruptive it's going to be. And it wouldn't be the first time that technology has come along and disrupted the job market, of course. And, and nor will it probably be the last, but what is it specifically this time?

Is it just the fact that, well, we need to, we need to be aware of what that disruption's gonna be and start managing for it, or do you think that there's a risk that disruption is misplaced in some way to your point? How sophisticated really is the technology and therefore, let's say there was a huge impact on.

Law for example, to pick a, to pick a relatively random but rule-based industry, do you think that putting something like a chat g p t into that too soon could be either something that we could then gonna need to roll back from? 'cause we find out, actually, you know what, it's not quite as good as we thought it was.

Or do you think that the more pernicious risk in that is how misleading a, a mis deployed chat G P T could be in important areas like medicine and law and other areas [00:07:00] of society? So let's take your last point first actually. So, law and medicine. So there was a recent announcement with Epic Healthcare, uh, and OpenAI that have, uh, that they want to actually roll out, uh, G P T and chat g p t into hospitals to allow people, allow surgeons to, to look up medical case histories, to look up at.

Look it up against patient medical. So basically e h R or electronic health records are all gonna be taken in and, and trained on, on, on the G P T and then the surgeons can use it to, to query and, and in the law there's already several cases where, or several examples



where people are using G P T to basically look through huge amount inordinate amounts of legal precedents and cases before to try and find correlation between the current case and another case.

Now the problem with all of this is that G p t, if it doesn't know it, it, it makes a convincing case for the answers that it presents. I mean, I've seen some of the [00:08:00] things And you would read it and you would go, if that came from a professor, I would take it verbatim. Yeah. But because it comes from a, a, a machine that you believe is infallible.

I'm going to absolutely take it as verbatim 'cause it knows what it's talking about and it's pulling from the internet and nothing, nothing is ever wrong of it with the internet and it's already done misdiagnosis. And a few examples that I've seen in healthcare and in law and in other examples where it's basically pulling in text references, it cites.

References that don't exist and it makes up case law and it makes up articles that doesn't, don't exist, et cetera, et cetera. So we have to be very careful of how much of the human in the loop we actually remove in this. So we're not gonna get to a state where we're fully automating some of this, but we are getting into a state where certain levels of jobs.

In these particular industries and white collar industries will be removed because of, you know, somebody's already talking about, well, I don't need paralegals because I can basically use this to do court [00:09:00] documentation, some basic research, you know, answer some queries, you know, and I'm a senior law, you know, senior partner or senior solicitor, you know, and, and I don't need to hire someone at 30, 40, \$50,000 or pounds or whatever denomination.

You know, to do it, I can just pay \$20 to a chat bot and to do it for me. And, and that's a really good point is that it, it's now challenging jobs that traditionally would never be under threat through the standard type of automation that we would expect. So we've been used to farming automation, we've been used to digital automation, to process records, et cetera.

Now we're talking about very high order jobs, having a material impact, and I think that's upset. Quite a few that didn't think that their job would ever be under threat when they entered into the industry. Um, and that's, that's changed the conversation quite a bit like the paralegal point, which is I can replace you for \$20 and good luck elsewhere.

Yeah. And I, you know, there's another example that I saw where, and, and you know, in, um, [00:10:00] biotechnology where someone is, is essentially using PhD level candidates or PhD level people to do data labeling. In bio labs as part of research, and he wrote an open sort of blog post that basically said, here's this person.

She has a PhD, she does data labeling, et cetera, and I don't see her having a job in six months. Now that's someone who's spent, I don't know, 5, 6, 7 years plus PhD on top, learning about biotechnology and bioengineering to have our first job on the ladder. Uh, doing data labeling and data sets for, you know, for, for research.

And that's gone. You know, and, and what I see, and it's very, it's interesting is I had another podcast with someone else, um, a friend of mine, um, last week. I. Is that the job ladder, in a sense, the first few rungs are gonna be completely whipped away. Yeah. Yeah. But then there's actually gonna be, and above, above the ones that are [00:11:00] left, there's gonna be another set that are gonna be completely whipped away because who needs managers?

What's the manager gonna be looking at? Well, the manager's gonna be managing what? Someone who's using an AI or a bunch of automated systems. Well, we don't need them. So there's actually gonna be several rungs missing in career ladders from going forward. Yeah.



And a big part of your, your career is learning on the job.

Yeah. You can come outta university and Great, I got my education, but without a, a number of years learning, you're not as effective as you should be with an industry. And how'd you do that now? If you're replaced by an ai? How'd you get to a high order? Capability, which an AI can't do and allow me to start a career path.

And that becomes tr, I mean, that, that could be a big issue about getting new talent into an industry. I see two dimensions to that too. So one is absolutely the point that you guys have just made, but the second one for me is then also what happens to that industry. So let's say you roll that clock forward 30 years, you've had, you know, kind of a, a, a number of.

[00:12:00] Experienced people retire from that industry who replaces them. Mm-hmm. And then, and who, let's say, take law as the example. Let's, let's say if you said, well, actually yeah, AI is going to replace a lot of the, the bottom rungs of the ladder that you set out, Theo. Um, but we're, we're always gonna use humans for the actual court cases in the decision making.

Well, that option removes itself. 30 years down the road if nobody's actually, you know, kind of in the pipeline to replace that thinking. Right? Yeah. I think, and, and, and you know, if you project forward 30 years, you've also got the, the, the 30 years advancement in, in AI and technology to the point where, you know, you probably won't need that level of experience anyway.

And, and the other thing as well that I, is, that I'm seeing or I will see is, Obviously you've got G P T as it as it stands just now, but you've got everybody building really cool stuff on top. Um, like auto, G P T and baby a g I, which are, um, AI agents. [00:13:00] So you can instruct G P T to go away and do something, and then it spawns almost, for want of a better word, subroutines, that will go away and complete mini tasks to draw it all in to complete the big task.

Um, and, and it makes it more efficient and, you know, it'll call on different types of software and stuff like that with different bits of data to an analysis to do the work and then draw it all back to the conclusion. Well, if you take, if you take that model, Trump three writes itself, doesn't it? Well, yeah.

Uh, but the other, but the really interesting thing, which I don't think anyone's really caught onto the fact yet, is that stack software vendors or enterprise software vendors are doomed. Why would I need to go and implement something as heavy and as Franken stack solution when all I need now is a a text box.

There's the fundamental point, isn't it that the interaction model paradigm completely changes, so that's more human-like, and then you get this, the way we [00:14:00] process that request fundamentally changes and you get a higher quality outcome. So the traditional transaction-based systems that we've been used to for decades, and they run global companies.

Yeah, absolutely. You can see the evolution away from that because you'll get better results faster with this style of technology. We're all gonna be moving backwards to essentially databases again, and it'll be the fastest, most efficient, most structured and unstructured databases. With the AI front end to be able to reach from it.

The beautiful thing about that is, is that the algorithm has fundamentally replaced the algorithm. So you've got this, you've got this thing about the algorithm that trains the data model that creates this outcome that we're talking about is replacing the traditional



algorithms that have survived for 20, 30 years on how we process these transactions.

So it at it's usurped itself, are you saying it's almost Darwinian? Yeah, it is. This is the next evolution of the algorithm which has killed off its old, um, its grandparents. Yeah. Yeah. Going in that, um, I, I wanna talk about [00:15:00] Greg Brockman's Ted Talk recently. I think, I think it only came out this week or last week depending on when we put this episode out.

And Greg, I think in his talk he demonstrated pulling together a number of elements of AI and automation. Uh, to create some new experience. Theo, do you wanna just maybe just set that out and just give us a, give us your take on that. 'cause I think this ties a little bit to your concern earlier in the conversation around automations as well as ai.

Yeah, sure. So during this TED Talk, and it was actually fascinating 'cause it was live, it wasn't recorded. So, I mean the guys got balls to be able to do this and pull it off and have the confidence kind of, sort of thing. But it shows to the power of, I think their a p l strategy. 'cause as we know, a couple of weeks ago they said, you know, we're basically opening up our APIs that allows everybody to hook into it, apply, you know, apply to use it and apply to be a company.

One of, uh, you know, 4,000. Well, actually no, it's one of a few, but one of them is Zapier. And [00:16:00] Zapier has 4,000 integrations already. So it allows you to basically, Who can different software applications as part of your workflow if you're building workflow. And what he demonstrated was, you know, I'm, I'm writing about something.

I want to, I want to cook a meal. This is the meal. And you know, it goes, you know, I want to cook such and such. Give him instructions to cook a meal. Give him the ingredients and remember the ingredients, which is a new thing that they're introducing, which is memory between chats. Um, It produced a picture of the finished product so you knew exactly what you were gonna get and your mouth would be salivating, et cetera.

And then it recalled the ingredients and went to Shopify, I think it was our Instacart, to basically find the ingredients, make the purchase, you know, and get it delivered. And that all happens from the one kind of sort of text prompt in a sense. But, and again, that shows the power of mm-hmm. Of, of these little agents that are doing separate things.

So I want to cook a meal, which shows I'm, you know, I want to have an intent to [00:17:00] actually do something and I don't have the ingredients and, you know, if, if I instruct it, so it's gonna go off and it's gonna do those things. It's really interesting because one open AI is watching what everybody's doing with its tools and then going, well, that's quite clever.

So we're gonna build it natively now. Yeah. So by the power of the development, open source development community, they're doing all the hard work and testing everything for them, and then they're going, oh, we'll cherry pick that and we'll cherry pick that and we'll put it into the next release. Suspect.

That's not the first time that's happened. No. In the tech industry? No. That was like a real company. Sorry about this. And then the sec, the second thing is quite interesting is that a lot of this stuff that they're doing now is what was, what was touted way, way back with I O T and with Siri and Alexa and everything else.

Yeah. But they just didn't have the technology in place at that point. But this is, this is a good point. 'cause if you take the traditional. Uh, bots assistance that we've had in the past, like I've just mentioned, the quality of the [00:18:00] output you get from this new generation. I get really frustrated now and go, ah, why can't you do better that can over there?

So they're almost like redundant. What came into our life five years ago is now going, well,



that's rubbish. It doesn't give what I need anymore. So my expectation has been lifted much higher with seeing the quality of the results that some of these engines can create. Yeah, so it's, it's now frustrating me to say, that's all legacy.

Move on, move on, get faster. In that interconnectivity that Greg demonstrated, I mean so powerfully actually in that, in that Ted talk that to. Feels like the sort of the, the second order risk. So the first order risk of what's going on at the moment is actually less about the whole singularity issue, and it's actually much more about unintentional and in your slightly clumsy disruption of professions in the job market.

Does, does the second risk come from the interoperability? Is that, uh, that Greg was talking about, do you think, and, and the automation that's, [00:19:00] that's then gonna be wrapped up in that? Yeah, I mean, if you look at. So we've all lived through, you know, our various stages of what automation looks like. Mm-hmm.

You know, we've had business process automation. We had case management, we had robotic process automation, all those kind of, sort of, you know, tools that we've had before. I mean, R P A was supposed to save everybody so much time and it ended up just being screen scraping data. Yeah. Whereas this is complete, you know, this makes r P a redundant.

Um, once, once. Once the next iteration, I think, you know, in the next five years, the level of sophistication between being able to instruct something to do something once and it goes away and completes the task in whoever manner, you know, you choose and then it comes in, you know, and it gives you the, the finished result.

So it could be go process this invoice, negotiate with their procurement bot and get me the best discount for that piece of software, and then come back and then raise the invoice, pay the invoice, blah, blah, blah. You could do that all in a string, you know, all in a sentence or a natural [00:20:00] conversation using natural language, and it would go away and do it.

And you might not even have any auth, auth, you know, authorization levels, you know? Yeah, exactly. In terms of, you know, uh, I can only spend five grand or whatever kind of sort of thing. Yeah. Yeah. Like you say, that second order, that next leap going on from what we're seeing here. 'cause it's all very consumer based.

That's right. You know, it's, it's like, you know, go and do my shopping for me. Go and ring up the hairdresser and book me a, an appointment. Yeah, that's, that's all very clever. And, and stuff like that. I want to see what people are doing with it in the back. You know, what is Microsoft? Telling everybody else, you know, on an enterprise level what it can do.

That's the problem. And also like any system that's invented for reasons like that there, there's, there'll be nefarious, uh, uses of it at some point as well. And I think you put that really well, which is like in the sort of, I dunno, just in that cooking a meal example. There were probably 10 or 12 different integrations going on there.

Maybe, maybe more, maybe a [00:21:00] couple less. None of which were making any judgment. Were bounded in any way, whether it be spend levels or, or much more complex ethical issues that might go on between, you know, the execution of one of those commands. Yeah, I mean the, take the shopping part, for example. I mean, I would love to know, you know, how much work you have to put in.

In order to get the ingredients you want because it could be, 'cause it could arbitrarily just go to the most expensive, you know, shop and go, I'm just gonna buy that. 'cause the availability's there. Yeah. And you've spent maybe double the amount that you want to do just for a single meal, or do you have to set a parameter and say, I only ever want to pay.



10 plus minus 10% variance on average shop prices for these items. And it might have to source 'em from five different things. In which case are you paying five different delivery charges and who you're buying it from and you know, do you have an account with it? And [00:22:00] it's, you know, there's all this stuff that we don't know yet.

And I think, I think, well, this is a, a really excellent point around passing authority to an automated system. Once you give an automated system authority and it executes at a pace we can't track from a human perspective, the risk of damaging unintended consequences starts to increase. Yeah. So it could wipe out a bank balance or something, or mm-hmm.

It What's the sustainability impact? Of it sourcing in a, in, in, in a bad way and things. And if you think about then that's applied to large consumer product organizations and they make sourcing decisions that, you know, are measured on a global scale, uh, you could end up getting some pretty bad unintended consequences off the back of that.

So what's the risk and how do we judge the risk associated with passing authority to the models? 'cause we've not done that yet. Mm-hmm. We're all looking at the results and then there's a human in the way to judge it. But at some point somebody's gonna take that leap. To the next stage, aren't they? And at that point then I think it gets very interesting.

Yeah. Um, I was with [00:23:00] defense organizations yesterday as part of a workshop on, uh, you know, various things and I was running the one that basically talked about generative AI and AI and things like that. It's very interesting to understand. Their risk appetite versus everybody else. Now you can, now, we're not talking about kill change stuff where it's like, just speak to the robot, and the robot goes off and rampage and kills everybody.

It, they're, they're even, they're looking at it more from an operational point of view, but they made a real, really good, strong point in that. Uh, unless you have the data in the right way and the right level of information available in your organization, this stuff is not gonna plug the gaps. Hmm. Um, so we do have an enough, you know, much in the same way that the early sort of AI.

Projects and the machine learning projects, everybody thought it was, oh, I'll buy ML and I'll just shove it in and it's gonna know everything. And it's like, no, you've got about two years to train the data before, you know, figure out the data and then train the models before it even becomes [00:24:00] useful.

Yeah. And you've still got that project and that lead time ahead of you, um, in order for a G P T model to make sense of anything in your organization and. We've got Bloomberg, G P T, who have basically, Bloomberg has basically released their own large language model, trained on their entire history of financial information.

Now that I don't know how long they've been at it for, but that to me says that the structure and the knowledge and the depth of work that they've done to make that data available means that they've done it from day one. Hmm Hmm. And this is what's missing in a lot of organizations is I've never had that discipline from day one.

Yeah. It's been a challenge for years now. Yeah, yeah. For many organizations. And that still remains. Yeah. Yeah. But if you, if you take the Bloomberg issue or the, what they've done, their business is built on analytics and structured data and being able to understand what's going on. So they probably started from a [00:25:00] very good place.

'cause you know, they're built, they have to have it. Whereas other organizations like say, shout, it's um, they've really struggled with data control. Yeah, I, I'd love to see what a bank does with this and how badly it. They, they mess it up. Someone's watch for in future. Um, so



I think that just then in, in summarizing what we think the, some of the ethical risks are, and obviously this is an incomplete list.

The first and most straightforward one is that it's not always right. And actually there are examples of it bringing in. Artificial information to try and make its own case and presenting that side by side with real information in a way that's indistinguishable without having it being declared as ai.

And we had Dave Snowden on and he made a great point about the fact that only three to 5% of human knowledge has actually been written down. So when you. Uh, in a position where your AI is actually using written down information, the issue is not AI becoming a threat to humans. It's about humans becoming as [00:26:00] dumb as ai, which that was a, an interesting way to hold that particular risk.

The second one, is it then in around the job market and profession disruption? In a way that might kind of harm what's going on in that profession, how people move through professions, and how you professionally develop yourselves though that being said, that's gonna run alongside the development of ai.

So there's gonna be a lot of iterations in that space that probably needs more attention than it's getting at the moment. And then the third big point we just talked through is when you bring in increased integration, a p l and automation. So you're actually setting off a series of linked. Events as a result of whatever command you made, like make me a bowl of tam tomato.

So the sort of ethics and barriers in that decision making is not necessarily sophisticated yet, and all of that before you even get to something like a g i coming onto the playing field, I think. That being said, then what are our tips for our thoughts on better managing this? So it's, it's [00:27:00] quite easy for the writers of the letters sit on the sideline and say more should be done.

But sort of what does that look like? And just as we were prepping the episode Q grant that well-known, that well-known observer and commenter on the tech world, as has, has put something out that I thought I'd bring to the show. So this is direct from Hugh. Dear tech scientists, I'm one of those people who are terrified of the advance of ai, but ultimately, if it became a problem, couldn't one just, um, unplug it at the wall?

A serious question. How can it stop? Um, so Theo, what, what's, what, what would be your response to Hugh? It's too late for that. Hugh, I'm sorry. That's the Rubicon. We've gone, you know, this is Pandora's box. It's not a simple case of pull out the batteries, um, or unplug and switch off. I mean, you've seen basically, you know, people are, are creating their own versions now I.

Um, they're grabbing data where they can, they're creating their own LLMs. They're training it, they're using open source models. Um, [00:28:00] you can buy Alpaca, which is the Stanford University version for about \$600. So if you've got, you've got \$600, you buy your own, host it on your own machine and start using it.

Where's the control? Where's the plug? You know, um, you know, that's local, that's a local machine. Funnily enough, the same Ted, um, had Humane's, uh, c e o Imran there, and they've been working on, uh, a secretive kind of, sort of AI hardware project. And it, it, it, to me, it's like a. It's like an extension of the Star Trek Communicator.

Badge, yeah. But with a, a projection so you can project images onto your hand or, or a surface and interact with it. Now, removing the, the stupid gimmick side of things, the actual



idea of having your own personal communicator or your own personal assistant trained on your ar on your data. Mm-hmm. Which is the, the real.

Critical point to make here enough people want, want to, will want to use that for whatever, you know, reasons that they want, you know, that, that, that they feel, [00:29:00] that sees that they see fit. And it's that, that's a worry because you can't switch that off to me. 'cause you'll, you know, if, if. 3 million people suddenly go, I want that.

They buy it. How are you gonna tell 3 million people who have spent all that money plugging all their personal information, literally their life data, to train this thing, to be as effective for them as possible? To sit, to turn around and go, oh, we're gonna turn it off next week. Because, you know, Italy doesn't like it then, uh, it's, uh, yeah, you, you know, you're gonna have 3 million, 3 million, very pissed off people.

Well, this is, this is the thing, isn't it? You, you've, you're assuming that everyone's. All well intended, but there's lots of nefarious states in the world that'll just continue to develop it and use it. And you can't absolutely control everybody all the time. It's already accessible and out and Yeah, absolutely right.

The box has been opened, we just go away for it. Which makes me think maybe Frank Herbert's view of the world in *Dune* is actually going to become true in the future as the, uh, the reset when ai uh, [00:30:00] Everything. So maybe Hugh Grant's just the first to call it out. I think, I think it's interesting when, when people go, oh, turn it off at the wall.

And it's like, well, which version are you turning off at the wall? Is it the, the Western version? Is it the US Silicon Valley version? Is it the China, is it China's version? Because we know China developing their own, but they're quite happily gonna give it to all their citizens because it's essentially gonna be, um, a control mechanism, uh, control mechanism.

And they're, you know, it's only gonna give them the answers that the state wants to give. And, and to your point, AI is going to be the thing that creates competitive advantage in the markets and in industries. Defense being a massive example of it, as well as sort of personal hyper-personalization.

Nobody wants to turn it off. 'cause if they do, somebody else is gonna take control. So yeah, we've just gotta try and. Grapple it, haven't we? And but regulation and legislation isn't catching up with this, as we've seen with so many things in the tech industry. Regulation, legislation can't deal with globalization.

And that's been going on for 40 years. So God help 'em [00:31:00] when they come to try and make a law about this. I mean, Italy great first in, but you sort of think maybe you've overshot the, the point you somewhat Yeah. Somewhat unsophisticated in the blunt instrument. Yeah, exactly. And but, but, right Rob, turn off thing.

I mean, uh, before you even get to, like, which version of AI do you turn off and say, what do you actually turn off? Off? Yeah. You know, it's like you see we should switch off the internet. 'cause that's actually what it would come down to. You know, like, you know, one, it's, it's not AI running as a separate machine in each country.

Is it? It's like, it's it's absolutely embedded in almost every layer at this point. So that, that I think is a, just is, is a non-starter. Without the, without the doomsday scenario, you switch, literally switch. Yeah. Um, personally, I'm gonna welcome our robot overlords and just stoke they're benevolent when they arrive.

'cause it's happening, isn't it? Oh, just look at your Twitter history. Yeah. They're gonna, they're gonna look back on this conversation rather than gonna have you marked out right



from the start. Yeah, exactly. [00:32:00] Yeah. I'm gonna be, yeah, they, they'll like me, but the rest of you know you're in trouble. Yeah. You were like, no, I was there.

I was there defending it back in the day. Yeah. I was, I was all for it. Thank you. Yeah, yeah, yeah. So, so the only thing that we can do, embrace it and make sure that the facts. Or straight that coming out of it. Right. I, I think it isn't, it's, it's a ma It is definitely a matter of getting it more and more.

Correct. Yeah. The other thing that's needed for me is it needs to be declared. So, you know, kind of AI developed material needs to be declared as being AI developed material. I saw, um, a thing the other day, I dunno whether you spotted this, anyone, but, um, there was some guy on Twitter that said that he'd been looking at playlists on Spotify.

And he'd been listening to a particular song and I think he had like a little thing going compiling playlists for him automatically. And he was listening to a particular song by Artist X and then cut to 20 minutes later, the same song came on by [00:33:00] another artist. Of a different name with a different piece of, you know, extremely poor cover out.

And then he then twigged this and did a bit more investigation. He said very quickly he found 10 or 12 or 13 versions of that like the same song. There had obviously been AI generated with AI generated cover art that, and he was speculating that maybe the streamers were doing this 'cause they then don't pay royalties on it.

So if they've got ai, you know, creating content, it's their content. People can use it however they want so they can flesh out their playlists with it and all kinds of different. I mean, there's a whole other showing kind of the impact on the arts from that perspective. Um, because I'm not necessarily even trying to open that box, but it seems to me that there should be a declaration around that.

Then it's like, find that that happens, but that maybe should be declared. And then of course there's the point on rules and regulations that Rob made, which is, I, [00:34:00] I think there's a. Conversation missing here because of a lack of sophistication and the distraction in our government and authority bodies of just not understanding this.

There is a whole conversation missing which to go full circle on the conversation. Maybe is. Maybe that's where that letter is coming from, which is a call for that conversation. Do you think they're, well actually Dave, you make a really good point. I'm gonna pick up on the the declaration side because as more businesses.

Look to adopt this, it's gonna be more customer facing. Uh, and the thing is, is that you will, it will get to a point probably within the next five years where you will be, you know, it will pass the customer service Turing test, if you wanna call it that, which is, how do I actually know I'm talking to a. A chatbot or not.

Um, I mean, that's, that's increasingly difficult, frankly, from some of the quality of customer service generally. That's a, it's a pretty, that's a pretty low bar, I'm gonna say. Well, yeah, but you know, at the end of the day, you might have to declare, you know, right up front [00:35:00] you are actually talking to an ai, you will be given advice by an ai.

Yeah. You know, if the next time you sit down with a mortgage advisor, for example, are you actually gonna be taking instructions from an ai, um, who's, you know, who passed the, you know, the C F A. You know, exam or something like that, you don't know. No. Exactly. Exactly. And that, that to me, because you, you just know that regulation isn't gonna keep up.

No. And, and the conversation isn't gonna keep up. So something else has to happen, I think, to, to make us aware of that. Oh, it, it, it, it will become so prevalent. We'll just rely on it and



we can't remove it. But there's, there's got to be a set of events that will occur where. Things that we don't want to happen will happen and we'll just have to use them as learning events.

And it's generally that, isn't it? You can't trust the legislator to keep up. So we're gonna have to have something that happens that then corrects our thinking or processes around it, and it feels like that's what's coming. We're waiting for the unintended consequence to hit then a retrospective, and then are we best change that over there.

[00:36:00] It's almost like, we'll, we'll patch it, uh, you know, close the stable door after the horse has bolted type thing, because it does feel like that's what's about to occur. We're just. Waiting for what other set of events? Well, we've seen that already. Um, so there's the, the, the Belgian guy who committed suicide after talking to a chat bot for six weeks.

And he was obviously deeply depressive and he just, and it got to the, you know, he got to the point where he was talking about climate change and he had a concerns and he was ob and he had, he suffered from mental illness, but he used the chatbot as a kind of, Therapy session in a, in a way. And the thing is, of course, a, a chat bot is not gonna, you know, uh, convey empathy or understand the relationship or anything else like that.

It's just taking an instruction and spitting something back. 'cause there were no guardrails or safe safeguards in place. It, he, he basically spiraled to the point where the chatbot essentially said, well, if you weren't here, that would remove the carbon footprint that [00:37:00] was impacting the environment. And he was like, okay, then.

And that was it, you know? And then you've got, there was another interesting article with, um, people who had formed relationships, deep, personal and, and, and in some cases sexual relationships with these, with a chat, a a specific instance, um, of chatbot, not chat, g p t, but another one, another large language model.

And the developers updated it, updated the software, and of course that broke the connection. The, the, the relationship that they had built over a, a large length of time to the point where it had imp they had implemented safeguards and it became really jarring 'cause it was like, oh, you and I were, were gonna go to bed tonight, blah, blah, blah.

And it's like, I don't wanna do that anymore. Dave, um, that was my how impression, by the way. That's what I'm assuming, rather than the host of the show. And that's, and it's like, oh no, I've, I've completely lost my, uh, my [00:38:00] girlfriend. Um, kind of sort of thing. And so how do you, you know, how do you manage that?

And this goes back to the ethical side, the unintended consequences. You know, the use case, you know, the use cases that humans always find themselves in. It doesn't matter what you build something for, there will be someone who uses it for something that you. Didn't intend.

So each week I will do some research on what's trending in tech, and this week I want to focus on free conversations that organizations need to have about ethics and ai. So with the current adoption of ai, concerns about AI and ethical violations have become common in many organizations. But it turns out that for [00:39:00] most of these organizations, it can be really tough to turn these concerns into actionable conversations.

So machine learning ethics and their points of intersection bring a lot of complexities. There are no quick fixes for these complexities and conversations around this can feel very abstract. So to implement effective AI risk mitigation strategies, companies should begin with a deep understanding of the problems they're trying to solve, and getting to the desired outcomes requires learning to talk about these issues differently.

So first companies must decide who needs to be part of these conversations. So, for



instance, a senior LE senior level working group that is responsible for driving AI ethics in your organization. And then this group should define their organization's ethical standards for ai. So identify the ethical risks for their industry or their organizations and where they stand on [00:40:00] them.

Identify the gaps between where they are now and what their standards call for, and lastly, understand the complex sources of the problems and put solutions in place. So really understand the bias in AI requires, for instance, talking about the various sources of the outputs. So not only the training data, but also how the inputs are weighted and where thresholds are set.

So a question for you, Theo. What do you think of these points and do you think this will help organizations? Absolutely. So to me, I, I think organizations are on their, their first footing. This is the beginning of their journey, um, in understanding what they're doing. I think we've seen, I think we, we kind of, sort of briefly touched on it, that a lot of these companies, you know, some of these companies have done machine learning projects, AI projects, data projects, and they might, may maybe, um, further along than other people to take advantage of.

Ai. But I think the, the whole question around ethical, [00:41:00] uh, ethical use of data and the protection of people, not only from, from a job perspective, but also from the end user and, and the the end customer as well. These are all considerations that I think, um, these organizations are gonna have to seriously sit down over the next, you know, year, I would say and, and, and ponder before they even start to look at any implementation strategy.

Yeah, I, I think organizations, if they're not already though, talking about it, are late to the game. Right. Which is. A com it, it, I think you said a mixed body of people talking about the implications of AI for, for a particular organization, looking at both the, you know, the human side of it, so like HR as well as the more sort of hardnosed business side of it.

But it feels to me that in the way that sort of broader society as, as let the rabbit out of the hat, you don't necessarily need to do that for your organization and you can at least try and. At least try and control it from a business risk perspective, [00:42:00] perhaps. I think it's just important that, you know, like your report, but also just in general, I think businesses are really championing it a bit to sort of understand how they can get into this.

Yeah. But they have to just take a step back. There's, there's room for experimentation, but I think there's also room to understand just the limitations as well. But I think over the next sort of five years, We're gonna see who's adopting it, who's not, and you, you'll probably see the gap in terms of, you know, that is gonna be wider than it ha has ever been.

Yeah. If someone says, oh, I use Pega and you don't. Mm-hmm. That's almost gonna feel like, you know, minuscule compared to, I'm using an n l a large language model and you're not, that golf is gonna be like this. I think in the Mele that is, The corporate world that fills most of our economy. The adoption of AI at the moment is limited, so it feels like those organizations have a little bit more time to get the ethics point, right?

Yeah. [00:43:00] There's lots of organizations doing funky things on the fringe, but we still look about how the core of the world operates. Yeah. It's still on that traditional style of approach that we were discussing earlier, so hopefully they can get a grip of it before they start to implement it on massive scale.

But to your point, I can replace an E R P platform with this thing over here if. That romps over the rise in three years, they might, they might take that before they've had the ethics



conversation about the impact on their organization or the decisions that that might make on their behalf. Yeah. And it also depends on the data that they have, right?

Mm-hmm. Yeah. Yeah. Plays to the point that, that we discussed earlier that a lot of, a lot of organizations still don't have a handle on the amount of data that they, they have in the organization, how it's structured. You know, there's, and actually you made a really good point earlier as well, about. Tacit versus explicit knowledge.

So explicit knowledge is obviously being recorded as data, but how much tacit knowledge in the organization exists that is not touched by data or, uh, cannot be touched by ai. Yeah. Um, and [00:44:00] that could be a saving grace for a lot of jobs, to be honest. 'cause how do you capture someone with 10, 15, 20 years worth of experience that knows the business, knows the customer base, you know, that an AI doesn't have?

And that's, that's a funny thing 'cause all those business developed apps that we've always wanted to get, get rid of 'cause they sat in Excel spreadsheets on file shares are actually going to be the savior of the employee in the end, isn't it? So it's actually, it's actually all those little apps that carry all the risks that, and nobody actually wants actually, yeah.

Save the jobs. You imagine how powerful AI would be though when it realized that and it goes around and gathers 'em all up. Just like a stock take of all the little spreadsheets and stuff like that, and suddenly you will not be able to get a car parking space without talking to the ai.

Thank you. She and everybody for I think. One of the most important conversations that we should be having as a human race at the moment without overstating the [00:45:00] grandness of it. Uh, obviously this is a tiny part of it, but I think we would encourage anybody at this point to be starting to think about this stuff and whether they're applying it in their home life or their business life to be taking it seriously, um, because it should, it should be taken seriously.

I, I think one of the things that the letter definitely does get right is about the profound impact it's gonna have on, uh, us as a human. Society. So Theo, uh, great to see you again, and we end every episode of the show by asking people what they're excited about doing next, which could be, thank God it's Friday, roll on half, five, all the way through to, uh, some exciting you're doing in your business life.

So, Theo, what you up to next? So, I've been stuck in doors actually, but, but, uh, Tending to an ill child for five days. So I'm actually looking forward to getting out and watching. Um, evil Dead Rise. Oh, which is completely random. I'm probably watching the finale of Picard maybe two or three times more. I think I'm gonna watch that tonight.

I'm, I'm actually looking [00:46:00] forward to that quite a bit. And Evil Dead Rise as well. I was, I listened to a podcast the other day, uh, called The Big Picture, which is a really, really good movie podcast, and, um, they really liked Evil Dead Rise. It seems to be getting well liked, isn't it? Cool. Good stuff. Yeah.

Good. Well to know. Enjoy, man. Enjoy. Thank you. So a huge thanks to our guest this week, Theo. Thank you so much for being on the show. Thanks to our producer Marcel, our sound and editing wizards, Ben and Louis, and of course, to all of our listeners.

We're on LinkedIn and X, Dave Chapman, Rob Kernahan, and Sjoukje Zaal. Feel free to follow or connect with us and please get in touch if you have any comments or ideas for the show. And of course, if you haven't already done that, rate and subscribe to our podcast.

See you in another reality next week [00:47:00]

About Capgemini

Capgemini is a global leader in partnering with companies to transform and manage their business by harnessing the power of technology. The Group is guided everyday by its purpose of unleashing human energy through technology for an inclusive and sustainable future. It is a responsible and diverse organization of over 360,000 team members in more than 50 countries. With its strong 55-year heritage and deep industry expertise, Capgemini is trusted by its clients to address the entire breadth of their business needs, from strategy and design to operations, fueled by the fast evolving and innovative world of cloud, data, AI, connectivity, software, digital engineering and platforms. The Group reported in 2022 global revenues of €22 billion.

Get The Future You Want | www.capgemini.com

