

TECHTalks | EPISODE 4 | GENERATIVE AI

With Ed Bowen, Managing Director, AI Center of Excellence in Deloitte leader for Advisory, US Deloitte Risk and Financial Advisory LLP, and Costi Perricos, Partner, Global AI and Data Leader, Deloitte UK, Global AI and Data Leader

Raquel Buscaino: Generative AI has advanced significantly in recent months. It's allowing people and organizations alike to generate text images, audio, and more from user-generated prompts, unlocking human creativity across industries.

It's democratizing AI applications and taking the world by storm. And it is just the beginning. Welcome to Deloitte TECHTalks. I'm your host Raquel Buscaino and I'm the emerging technology lead on Deloitte's US Novel and Exponential Technologies Team.

My team senses and makes sense of emerging tech and on today's episode, we'll explore generative AI. I'll be speaking with two Deloitte AI leaders; Ed Bowen, Deloitte advisory AI leader and managing director of the AI center of excellence, and Costi Perricos from the UK, and Global AI and data leader.

Ed and Costi welcome to the podcast, it's so great to have you. Ed, why don't we start with you? Can you share a little bit about who you are and what you're doing in the space?

Ed Bowen: Yeah, thanks Raquel. yeah. Thrilled to be here. So, I joined Deloitte about 3 years ago, coming out of 20 years in life sciences. I lead an AI center of excellence, and really everything that we do in AI across Deloitte Advisory. So, our mission is really focused on risk. So cyber risk, fraud waste and abuse, anti-money laundering to carrying internal controls.

Raquel Buscaino: That's great and Costi, how about you?

Costi Perricos: Raquel, I'm really thrilled to be here, and thanks for inviting me. So, my journey with AI started about 25 years ago, when I did a PHD in the subject, and we kind of knew then that it was going to change the world, we just didn't think it was going to take 25 years to do so. I have been with Deloitte for about 20 years, I have been working in data throughout all this time. Be it content management in the nineties, and of course, in the last sort of 6 or 7 years, refocusing on my passion which is AI.

Raquel Buscaino: Well, truly thank you both for taking the time to be here, and I'd like to just dive right in and start talking about Generative AI, so, Ed, I might start with you and ask what is Generative AI? Can you give us a high level for those that might not be aware of it?

Ed Bowen: Sure, Raquel, there's certainly a lot of information about it on the web right now, but to summarize, I would say it's the ability of machine learning and artificial intelligence to create new and novel text, images, video: all kinds of things. I think we're just really starting to scratch the surface, using traditional techniques.

So the power of computing and the power of being able to aggregate large, large, large amounts of data, so imagine everything on the Internet, everything that's ever been published, and pull that together and build very complex models with trillions of parameters have really superpowered this idea of using AI to be creative.

Raquel Buscaino: And so you mentioned a couple of things there about all the different parameters. I mean, how does generative AI actually work? And what makes it different than the AI we might have been used to, maybe even a couple of years ago.

Ed Bowen: Yeah. So fundamentally, AI models and machine learning models, which are one of the same, are used to make predictions. And so those predictions might be something as simple as trying to tell the difference between a cat and a dog in an image.

Or something really sophisticated, like trying to predict the next word that should come in an article that's being written, or some type of a reaction to a prompt, and, with the large language models that have been incorporated into these generative AI tools, they have become very, very sophisticated at using the large corpus of knowledge around how humans write and then use that to predict the next word in a sentence, and then just keep building on that.

And so, it gives the appearance to us as humans that the output is something that a human might write, but it's new and novel. They've also done some really clever things like, add a little bit of randomness in predicting what the next word is going to be, to make sure that every response to a prompt isn't exactly the same and then by inserting that randomness, the actual word that follows that one that's just been predicted might be slightly different, and so you can end up with some really interesting outputs, but it also adds in into the risks which I think we're going to talk about a little later.

Raquel Buscaino: And so, Costi and your perspective. What makes the time we're in right now such a pivotal point in AI's trajectory?

Costi Perricos: Well, that's a really interesting question Raquel. I mean, if you look at the history of Generative AI, probably the biggest pivot point was an article published by Google in 2017, called "[Attention is all you need](#)" where the transformer model that Ed referred to was introduced, and that really kind of made the prediction of the next word, that was described earlier, much much more efficient, much more accurate. And that technology has been around, you know, since then.

I think what made the big difference was when OpenAI made this technology available to everyone through ChatGPT and all of a sudden, it has sort of democratized access to this type of technology. And, of course, you had millions of humans look at it, and you know, us humans are wonderful, wonderful beings which have imagination, we can think of all sorts of uses for technology like this. And it just took off, which is why so many people these days are talking about it.

Ed Bowen: I think, adding to what Costi said, I think there was a barrier for application around how accurate the outputs were, or how trustworthy the outputs were. I think it was pretty easy, you know those first transformer models when you would use them, some of the outputs, you would read it, and you would say: "Well, that doesn't really make sense, or that's not what a human would write."

I think, with ChatGPT, and GPT3 as the underlying model that feeds ChatGPT, it was the first time where laypeople, reporters, other people who are interacting with the technology were blown away at the power and the relatively human-like responses that we're coming out of it.

These models are likely going to continue improving, and it's that sort of performance barrier that's been crossed that's going to drive wide adoption, but it's also going to introduce a lot of risk in the use of the tool, because it's going to be so good that laypeople will read it and assume that the output is a 100% factually

accurate. And we know that there are still some issues in there around the models, doing things like incorporating information in there that's not actually accurate.

Raquel Buscaino: And so, both of you just touched on, I think, is that we're seeing this technological progress, but there is a bigger question, too, which is almost the human element of trust is, do I trust this AI output?

Can you speak a little bit though, to some of the barriers on that trust lens. If people are going to effectively use AI, what's the level of trust that we need? And how do we create that trust in AI programs, AI companies, applications, the rest?

Costi Perricos: To me, that is a huge point which I think both organizations and individuals are struggling with. As Ed mentioned you know coherency and accuracy are two different things. These models sound very coherent you know, they're 100% confident and coherent in the outputs, but they're only sort of 80% or 70% (*number estimates provided as example*) right. Now, they will improve over time, but you cannot take as face-value what these models are saying.

The way that I see it, is this almost sort of a two-by-two square, where on one axis you have, how long it takes to produce content and on the other you have how long it takes to review content and if there is a use case where producing content takes a long time, and reviewing it takes a short time: that is a great case for generative AI, because you'll have a human who can review it, you know, an experienced human who can say that's correct or not.

Where actually it doesn't take that long to produce content, but actually verifying it takes a long time. That's probably not a great use for Generative AI.

The other aspect of trust is trust around data and privacy of data. And I think that's one of the big issues that many businesses are struggling with. You know, we at Deloitte, we run a number of experiments with these large language models, and we found that the output that you get when you have some curated data, i.e. your own data that you have amassed rather than that the data that the large language models are trained on the Internet. The output that you get is orders of magnitude better. Which actually is great because it means businesses can differentiate themselves that they're not all getting the same, the same types of these language models.

The problem is: how do you safeguard the data that might be private to you as an individual or private to you as a business, and make sure that the language models don't get trained on them, and then get that IP sort of get shared with everyone?

And again, this is a very fast developing field. We know that all the companies that have got LLMs (large language models) are looking into that, but at the moment for what we are seeing, it's a real challenge in terms of really using LLMs at scale within businesses.

Raquel Buscaino: Yeah. And beyond that just the privacy. I mean, do you think that's where the trajectory is heading where we'll see these LLMs, these large language models, being fine-tuned for specific application areas. I mean, it seems like every other day you hear about a new company that's being spun up where it's using a foundation model and fine-tuning it for a specific use case in the legal or the life sciences health care or the industrials sector.

Do we think that trend will continue where you'll see a few big players with these foundational models, but then a bunch of smaller players fine-tuning it for specific industry applications?

Costi Perricos: I absolutely do. I mean, I think the sort of industry-specific use cases are probably the least talked about but the most powerful. We know that this kind of technology is being used for things like drug discovery. It's even being used to make advances in things like nuclear fusion. And I think that's where the real excitement is rather than just perhaps the content and text generation use cases we hear about.

Raquel Buscaino: Yeah. And Ed maybe we can dive into some of what those use cases are. What are some of the not-being talked about application areas that people really should be talking about?

Ed Bowen: Yeah, I mean. So you know, I agree with Costi, in terms of the industry specific applications, and I don't think it's going to be limited to just large language models. There's going to be other types of foundational models out there, and by foundational, we mean, aggregating large amounts of data and having significantly expensive computing costs associated with it to build the model.

I mean, prior to OpenAI's release, DeepMind has their protein folding model, called AlphaFold, that is used for predicting how proteins fold, which is a very important step in drug discovery, is understanding how a drug target protein will fold, so you can figure out how to bind a molecule to it and regulate the disease.

You can imagine things like a biological network foundational model. You know massive graphs that represent how protein interactions happen in a human body. So you're trying to understand how dysregulation of proteins or genetic variation may affect protein interactions.

And if you can model that disease correctly, then you can likely figure out, where your drug targets are, and design a molecule to be able to regulate that disease. Certainly you know those are health care examples.

We're also seeing a lot of other types of examples like in cyber security where you can take large language models and be able to incorporate them in/detection type of models and algorithms, and be able to give context to what you're seeing in terms of an adversarial attack on a network, or, any kind of lateral movement, or other types of attack vectors that companies are traditionally trying to defend against adversaries.

We see it in areas like fraud, waste and abuse, where you can provide context around fraud signals and alerts and use the models that way. So, I think it's both going to be applications of these large language models potentially fine-tuned and then also, you know, native new foundational models in whole new domains that have a completely different set of information in them.

Raquel Buscaino: I love those examples because I mean, we're talking about protecting human beings. Right? Two very important areas here, and it goes beyond just simply having an AI generated profile picture or having a better Chatbot but Costi, what about you? What examples do you see here?

Costi Perricos: Well, I think I think one of the areas that that is a little bit understated around what these models can do is in the areas of software development and coding.

What we are seeing, what a lot of people are looking at is how you can generate code. So, you know, in the same way that 20 years ago we thought languages like Pascal and Basic were great because you didn't have to program and assemble, and you could program it something that sounded like English. Now you can actually program, you can get these language models to create proper code by just inputting proper English. That's one area.

I think the other area that that is really important is the whole code documentation space. You know, coders generally don't like to spend time documenting code.

And, you know, particularly in areas like financial services, you have huge amounts of legacy code written in languages like Cobol on critical systems, where the people who wrote it are either no longer with us or have retired.

And so, this can solve some of the huge challenges that the software industry has around being able to convert legacy code into new code. Pretty much automatically.

Raquel Buscaino: I think the area to your point of software development code generation is so fascinating. And what I like about it, too, is it's about reducing the friction of creation, right? It's that I, as an individual, don't need to know as many fine tuned skills as I once previously did, to have to be able to turn my ideas into execution and that, I think, unlocks so much human potential.

Costi Perricos: Absolutely.

Raquel Buscaino: There are certainly risks that we should be talking about as technologists in this area but we also have to recognize that this fear of AI replacing us has been around for so so long, and it truly is an [Age of With™](#) and to your point, there will always be a role for humans in the future that we're building.

Ed Bowen: And I think that role actually is a liberated role. Right? So, I mean, imagine, if your job is to go in and do some repetitive tasks and you get bored with it, or you know it's not very fulfilling. Now imagine that we're able to go in and say, "okay, well, we're gonna have the generative AI create the majority of what you used to create, your job now is to review that, make sure that what's in there is accurate as a subject matter expert, and then innovate on top of that"

Right? I mean the innovation possibilities and we talk about work right now and industries, but it's also going to be into your private time, and you, you know things, you're passionate about, and so you know, we used to say this in Life Sciences, and now that we still do is that, you know, we don't want our scientists spending time moving data around or doing some type of manual tasks. We want our scientists designing experiments, we want them interpreting experiments, we want them doing science. And so, if we can liberate people's time to focus on the things that they truly add value in, then you know these types of creative AI, this type of foundational AI, can just be a huge enabler for humans.

Costi Perricos: I think that's right. Although I do think that fundamentally we're going to have to change how we learn as humans.

Because if you come to think about it, a lot of the learning processes are about aggregation of content, and therefore absorbing and learning that content, while you aggregate it. You know, be it, at school, having to write an essay on Benjamin Franklin, and then going on the Internet and researching Benjamin Franklin, and coming up with all these facts.

Well, now, you can have an algorithm that creates all those facts and the learning will be around discussing, maybe "Benjamin Franklin's effect on society" as an example, and you know it's the same in business, in many businesses, like ours or other professions that rely on insights.

When you first join a business, you may spend a lot of time aggregating information, sort of pulling together reports, and having the more senior experienced people look at them and derive insight or basically drive value out of these things.

If the task of producing these reports is automated, right, how do more Junior folks in companies learn? I think it's an interesting conundrum that we will have to solve.

Raquel Buscaino: Yeah, I think it's both on the lens of changing what we learn, as you mentioned, you know not just the facts, but the impact, the application of what we learned, but also I'm really excited about how we learn too. I'm actually really excited for a future in which hopefully, with AI assistance, we can start teaching people, in the native way that they prefer to learn.

Ed Bowen: Yeah. And I think a lot of organizations as they adopt Generative AI, they'll, you know, it's kind of like how technology rolled out. I expect that there will be a lot of "how can we incrementally improve a process by putting AI into the process or generative AI into the process?" and we'll be able to continue to improve productivity, and all those kinds of things that historically have been measured and managed.

And then you're going to have these innovators who are just going to go "Well, instead of trying to, incrementally make each part of my process better, what happens if I totally reinvent my process from front to back?" And just say: "I'm going to build this with AI in mind". As you do it that way, that's where you're really going to see these kind of "aha" moments.

And I think you're right, Raquel. It'll likely be in areas of education. It's going to be in areas that right now we're not even thinking about. There's going to be these, fundamental shifts in the way that core areas like education get managed.

Imagine an AI assistant that's working with a child and a teacher, and it's adapting based on how the child is performing in basic things like reading, and writing, and maths, and science. And the next problem that it gives the child is not some pre-determined maths problem, it's actually doing an assessment as it gives the child the problems and it focuses on the areas that the child is struggling on.

So, you can see all kinds of situations where these kinds of concepts are going to individualize performance improvement, and learning. So, imagine people mid-career having things tailored to them that are going to be sort of just-in-time, based on the role that they're in, based on where they aspire to go in their career, based on the things that they're already good at. It's a pretty exciting future that we're headed towards.

Raquel Buscaino: And I think you know the future of how generative AI as a whole is developing, you know, we hear "prompt engineering" so much these days. "Prompt engineering is the career of the future". "Everyone needs to be thinking about prompt engineering". In my mind, I think that's just the next iteration of how we access AI, and I actually think that as these generative AI foundation models start improving, that actually, it's going to be able to understand us humans more intuitively.

So I would love to hear your thoughts on this prompt engineering. Is it the career of the future? How does it evolve over time? What are your thoughts there?

Costi Perricos: That's a hard one. I mean it's interesting because that you know my son is now going through college, and about 10 years ago I told him that he had to you should become a data scientist because they were going to be the new rock stars.

And then a few years ago I told him he actually should become a machine learning engineer because they were going to be the new rock stars, and now everyone's saying that prompt engineers are going to be new rock stars. But I'm with you. I think prompt engineering is something that will be very, very important, particularly in the early days of these large language and these foundational models.

But over time I think they're going to become smarter. I think reinforcement learning will become more automated, in the same way that we saw big advances in what Deep Mind was doing with 'Go' when it started learning off games that he was playing with itself. So, I think all these things will become smarter and much more efficient, as time goes.

I think you know, to your point, I wonder if there's going to be a job, in a few years called an "Explainability Engineer", which is actually not someone who prompts these models, but actually can understand how they work, and where the outputs that that are generated, explain where they're coming from, and how they're being generated, because to me that is actually probably a bigger challenge than prompting.

Ed Bowen: I think Costi is right, it is fascinating to think about these things. It's amazing to think about the careers that don't exist today that our children and our grandchildren will be working in when they are adults.

I think that prompt engineering will probably be something akin to a skill set that a lot of people will acquire and use in their work, and it'll be about learning, you know, how to be skilled at, how do you get the most out of generative AI models? And you know right now, we're focused on large language models for, you know, obvious reasons. There's been some big advances in the last few months. I do think it's going to extend, you know, beyond large language models. So, it's going to go into other areas around images and video.

And so how do you get the most value out of these models? How do you coach the model to give you the best output for the purpose that you're training it for. I mean already we see with GPT4 that you can put images and text into that model to get output. So, the idea of multimodal inputs, how do you engineer those things to be able to get the best output is definitely going to be a learned skill, and I think a taught skill, right? I think that it's something that we (*generically*) will be teaching people how to do.

Raquel Buscaino: Yeah. And when it comes to coaching these models, as you mentioned, I mean one of the big conversations around AI, not just meaning generative AI, but AI in general is around bias, right? It's. How do we create AI systems that work for us? Are you seeing the emergence of AIs that are coaching other AIs on how to remove bias?

Ed Bowen: Yeah, so I think, just, my fundamental response when we start talking about bias and every model is biased there's always going to be some bias, because there's some data or some event that the model hasn't seen that could occur, right. So, you've always got some bias and a lot of times the bias comes because we collect data for some other purpose.

So, in health care we see it a lot where health care claims are processed for payments, is in an invaluable resource to understanding disease prevalence, being able to understand disease progression: How long does it take for diabetes to progress to diabetic neuropathy? You know those kinds of questions.

But the data was not collected for the purpose of research. It was collected to pay claims. And so that data over represents relatively healthy people in the population. Right? The people who are lower socioeconomically maybe can't afford insurance. And so, we don't have their claims. So, you know, every single data set you use is biased. So, I think it's always about understanding the bias and controlling for the bias right?

You need to understand it. What is the nature of the bias, and you know, both in a traditional AI model that we've been working with, or in these generative AI models. And I think, having mitigations in place is hugely important. And not just one mitigation, right? So, this idea of layers of mitigation. You're probably going to

deploy generative AI in the context of some more complex set of components that are going to give you some kind of business outcome.

And so, having mitigations along the way to be able to check that the model is giving you an answer that has some fundamental bias, and most importantly, if the model is supporting making a decision, that those decisions are fair and equal across different classes of people, whether it's ethnicity or gender, but you know that we have that kind of fairness in our technology is incredibly important.

Raquel Buscaino: Ed and Costi, Wow! This has been so phenomenal! As we think about wrapping up here. Do you have any final thoughts that you'd want to share with our listeners as they think about how they're approaching their own AI journey, or just generative AI as a topic in general?

Ed Bowen: Yeah, I think, having a plan for how you're going to manage risk with generative AI becomes really important. Things like, when you put generative AI models out there, maybe they're facing your customers, maybe they're supporting your internal teams. You're not going to want to just deploy those things and leave them out there running on their own, right? You need to have checks in place, you need to have humans in the loop, and you need to have an approach for that, to monitor that the AI that you're putting out there continues to perform as expected. So, I think, that becomes a really important part of the implementation.

The only other thing I would add is, this environment is evolving so rapidly. Things are changing very, very quickly. So, making sure that you, as an organization, you're staying aware of what's happening in the environment, and how that might be affecting your business, how it might be affecting your risk profile, what your competitors are doing in this space, how they're interacting with your customers. I mean, all of that becomes really really important. So, the sort of generative AI ecosystem sensing becomes important, and the risk management obviously becomes hugely important.

Raquel Buscaino: That makes sense. Costi, what about you?

Costi Perricos: Yeah, I think Ed is spot on, particularly on the point around companies needing to keep up to date. And I think you know, never has the sort of paradigm of failing fast, failing cheap and failing safely been so important. You know, companies want you try things out, and show that they kind of you know, understand the value that's being driven, but if it doesn't work kind of move on, and keep up with the developments.

I think for me, you know, the closing thought is that I'm actually really excited by where we are. We talked a lot about the risks. But for me this really epitomizes what we at Deloitte call [The Age of With™](#), you know: Humans with machines have proven to be better than either humans alone or machines alone, and I think if we can learn to harness this technology safely, it's going to take both productivity, but I think also humanity to a different level.

Raquel Buscaino: I love that perspective!

Well, Ed and Costi, thank you so much for shedding some light on the topic of Generative AI, from its accelerated adoption to its cross-industry used cases, I truly do think that we are going to see the technology impact the way we work and live in such profound ways.

And to all our tech savvy listeners out there. Thank you for tuning in until next time. Say Savvy.

Legal information

This podcast contains general information only and Deloitte is not, by means of this podcast, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This podcast is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional adviser. Deloitte shall not be responsible for any loss sustained by any person who relies on this podcast.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee (“DTTL”), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as “Deloitte Global”) does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the “Deloitte” name in the United States, and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.

Copyright © 2023 Deloitte Development LLC. All rights reserved.