Kriti Khullar:

Hello and welcome to the Actuaries Institute podcast. My name is Kriti Khullar and an Associate of the Institute. I'm currently a lead at IAG, and today I'll be discussing education insights on generative AI.

I'm joined by Amanda Aitken, Chief Examiner for the Fellowship Level Data Science Applications Subjects, Andres Villegas, Senior Lecturer at University of New South Wales School of Risk and Actuarial, and Jon Shen, Head of Solutions Engineering at Arturo and the Independent Examiner for Data Science Applications.

Hello and welcome.

So, over the past six to 12 months, generative AI has become a frequent topic of conversation. Andres, can you please provide a quick overview of what generative AI is for our listeners?

Andres Villegas:

Yeah, thanks Kriti. I guess everyone probably has heard of ChatGPT, which is basically the most common tool that we know in the generative AI space, but these are broader things.

So, when we say generative AI, basically what we're doing is using algorithms so that we can take an input to generate another output. So, it could be input text to generate more text, but in general it could be input text to generate an image or input an image to generate another image or input text to generate sound, video.

So, all these algorithms transform an input into another output that it could be text, audio, images, sound. That's what we broadly talk about, generative AI.

Kriti Khullar:

Thank-you. Jon, I might through it you way next. What was your first personal experience with gen AI?

Jon Shen:

Thanks, Kriti.

Back in 2019, there was this website I found, which was called Talk to Transformer. And the way that it worked was you were able to type in some text. For example, it could be the first sentence of a story like, "There are two people walking in the jungle." And then what would happen is the language model would then take that text and it was GPT-2 in this specific case and it would extend that text. It would take that text as initial starting point and complete the story for you.

Now, at that time, the quality of the outputs being produced by the models were pretty average, but even then, you could get an early sense that there's a lot of potential for text generation and generation of other types of content. And also, thinking about how it could be applied in business. So, for example, even at that time, I get an early sense, there are ways of generating text to, for example, summarise the key observations from a graph or maybe automatically produce parts of a business report.

So, I think no one really anticipated just how far the technology would progress in the next three years.

Kriti Khullar:

Thanks Jon. Amanda, what about your experience?

Amanda Aitken:

Well, my first and probably second experience both linked to Andres and Jon interestingly.

So, I first heard about ChatGPT becoming available in a Data Science Education faculty meeting. I think Andres, you're the one that raised that. And I thought, "Wow, that sounds interesting." And we were just talking about ways that it might be able to help students and also, ways that we might have to look out for students potentially using it to say, answer an exam question, which makes it harder to assess whether they really know what they've learned.

So that was my very first awareness of tools like ChatGPT and other gen AI tools coming onto the market.

But the first time I really saw it in use and particularly related to my work was when Jon fed through I think a past exam question for the DSA subject into ChatGPT and asked it to answer it. Some of the answers were really impressive and then some of them fell a little bit short, but I think I was particularly impressed with the way that it could integrate the context of the questions into the answer.

And I guess I just wasn't expecting it to be quite so advanced.

Kriti Khullar:

That's quite innovative for that point in time. Andres, may I please hear about your first personal experience?

Andres Villegas:

Yeah, so for me, interestingly the first one I had interaction with was generating images. So last year there was a lot of buzz in Twitter, when Twitter was a thing, in March and midyear about DALL-E, which is like a play word for DALL-E, but it's like with WALL-E, what you do is you can input text and generate images.

And I had followed it, but sometime in September, they released it for users, so they made it publicly available. I remember that weekend I just went crazy and obsessed and just inputting things and generating fun things and tried to see and it was impressive.

And my wife happens to be in the artists world, and it was Sculptures by the Sea and we always joke about the descriptions don't make any sense of what they have of the thing. So, we started playing like, "Let's put the description and see whether whatever it creates looks like the actual sculpture," and some of the results were quite impressive.

And after that I just started following and then ChatGPT came and GPT-3, that has been for a while, but then the obsession just stayed and then became bigger with ChatGPT in December.

Kriti Khullar:

Wow, Andres. Good application at Sculptures by the Sea. Jon, I'll pivot back to yourself for your perspective as an educator?

Jon Shen:

Well, exactly going back to the example Amanda raised earlier, when ChatGPT was first released last November or early December when we started to really play with it, I started seeing all these posts flooding my LinkedIn feed.

Everyone was going crazy about ChatGPT and obviously like as mentioned, I grabbed the exam from the previous month that had just been sat by students, put the questions in and question one I thought, "Okay, we've copied the question from the exam paper into ChatGPT, it's giving a reasonable response. Maybe it's just one question, it could just be lucky."

But then putting in question two, three and four and then with each additional question being more and more amazed by how legitimate the answers looked, you've got a real sense that it wasn't just hype. There's something here that we should be very seriously thinking about as educators because obviously if a student can potentially pass the exam by spending five minutes copying and pasting responses from ChatGPT into the exam script or as educators, we really had to think deeply about changing the way that we assess student understanding and finding new ways to assess whether the students are learning.

So, that was one thing for us to consider. How do we change the way that we assess students to understand what they're doing? And at the same time recognising that it creates opportunities for us as educators.

Firstly, being able to brainstorm, looking at new case studies, new scenarios that we can use as learning exercises. And secondly, also having the potential to ask deeper questions. Because if you assume for a moment all students have access to ChatGPT, everyone is on that same level playing field, then it becomes less valuable and less meaningful for us to ask lower level or basic questions on what's the definition of X?

We can start to shift our thinking and focus efforts on what are the higher order questions that really get students to challenge themselves and to exercise critical thinking? So, how do we start to shift the assessments and thinking into that direction?

Kriti Khullar:

Thanks Jon. It's great insights over there. Amanda, may I please pass it on to yourself on how it's impacted your practice?

Amanda Aitken:

Yeah, so in the eight or nine months that we've had tools like ChatGPT available, and I have to admit I've mostly used text generation tools as opposed to image or... Sorry, also code I guess, maybe that sits within text generation.

But yeah, I've found it really helpful in my role as an educator in things like I often get fantastic questions from my students and they're so fantastic that I think, "I'm not really sure of the answer to that." And so tools like ChatGPT have been great because it gives me a starting point to develop my thinking about, okay, what is an appropriate response to this question and how might others answer this question?

Andres Villegas:

So, this has has made us think a lot about what we do at universities as educators because it's changing the landscape of how people learn, how we are going to work and in general, we are training the people that are going to be using these tools later on. So, in my practice, the first thing that happened is in December we immediately tried to, okay, let's look and try to see what's going to happen. And what we did, we had a workshop with all of my colleagues in the School of Risk and Acturial and we just have them play for half an hour like "How is this going to impact you?" And many of them hadn't seen it and I could see at the faces of them they were like, "Ooh, this is possible." But then the idea was to get ideas of how it will impact everyone's practice.

So, we discussed things about the one that comes to everyone's minds, assessment, which probably come back to that, but also what are the skills that we're going to teach? How do we need to redesign our curriculum? And then the other side, how do we do some of the administrative things that we do as educators?

Kriti Khullar:

Thank you, Andres. I might stay with you for the next question. Is generative AI a tool that is damaging students or helping them to learn? Which way are we tipping the scale?

Andres Villegas:

I'm going to be optimistic and I will say it's going to help them to learn, but to get there we need to have a guided use of it. So if we help our students navigate it the same way as they help us navigate how we can use these things, we probably can improve. But if we just leave it at the while without any guidance, just use it, it can be on the other side because it might be that some of these basic skills that we want them to learn, it will be that they don't learn them.

But if we train them on what is possible and what is not possible, I think we can ask them to do things that we thought that couldn't be done. And as Jon was saying, like high level assignments and we can expect more of what they can deliver. So I will be on the positive side, it's going to help them there.

Kriti Khullar:

Amanda, what are your thoughts?

Amanda Aitken:

I think I agree with Andres. I think it's more helpful than damaging, at least at the moment.

I guess my only concern is I don't really know about the long-term impacts on everyone, including students and all sorts of learners. And this question did make me think about someone that I used to work for, Len Boem, who wrote a paper about the risk of outsourcing your brain. And I just can't help but think back to that, that there is this risk that we might outsource too much of the critical thinking, society more generally, but especially students, you want to make sure you can still be a critical thinker. And so I guess there's that risk that maybe humans will stop doing that higher level thinking if it's easier for a machine to do it for us.

Kriti Khullar:

Yes, this is also important in the growth of our careers. Jon, over to you.

Jon Shen:

At the end of the day, ChatGPT, gen AI, large language models, they're all just tools and it comes down to how do we use the tools? Do we use them in a way that's beneficial to us? Do we use them to augment ourselves or do we just use the tool blindly on its own and forget about everything else? So it does come back to that critical thinking piece, and if students use generative AI as a tool to support their learning, to help them probe deeply on this topic, to get a better understanding of material, then I can see it as being hugely beneficial for learning.

But if you just take it blindly, if we go back to the example that Amanda made about using it for coding, there are lots of anecdotal feedback from coders saying, "If I use ChatGPT alongside myself or GitHub Copilot or whichever large language model for that application, I can be getting between 50% and 100% improvement in my code speed." That sounds brilliant at face value, but at the same time, that curve might not be efficient.

It might have security flaws which are problematic as you start to release them into the world, you put them into production systems. It might produce answers, but answers that are wrong. So again, we have to think carefully about how do we use these tools.

Kriti Khullar:

Thanks Jon. Amanda, do you think education departments should be looking to employ fine-tuned chatbots to help students?

Amanda Aitken:

I think this is a really interesting question. I remember listening to another podcast that was talking about use of ChatBots, for example, to fill in for psychologists. And some of the advantages being that if someone's having some real issues and they need to talk to someone, maybe it's the middle of the night, maybe there's no friends that are available to call, obviously their psychologist's office will be closed.

So, in those sorts of situations, like in education, it could be really helpful to have someone to turn to when you can't access your chief examiner or your lecturer.

But then on the flip side, I just have a question about how do we make sure that the advice that's given by these chatbots is actually appropriate and correct? And who's accountable if it's not?

If your education department is developing these chatbots and releasing them under their banner, then what happens if a student answers an exam question based on advice from the ChatBot and that advice is wrong? And so, their answer is wrong, should they get some marks for that? Is the onus on the education provider to make sure that everything that the ChatBot provides as advice is correct.

Kriti Khullar:

Jon, I might throw it back to you. What are your thoughts on employing ChatBots to help students?

Jon Shen:

There is a big opportunity to have these tutors that are infinitely scalable, that are always there regardless of time of day, and also have the opportunity and ability to personalise their responses depending on what students are asking and the student needs.

So, I think there is a huge opportunity there, but exactly as Amanda has said, there are these risks around hallucinations or what's giving the wrong answers. And at this current point in time, it's not so clear as to whether there are ways to resolve that at all.

Of course, there are ways we can try and put guardrails in place, but it's not so clear whether there's a way to completely with 100% certainty say, "Yes, this bot is definitely going to give us the answers we want every single time."

Kriti Khullar:

Thanks, Jon. Andres, what about you?

Andres Villegas:

So, my colleague, Patrick Laub, who teaches the statistical learning class with me and the deep learning class on his own, he started trialing in our class this term with something called the 'Class Chat'.

So, the idea is it's a website that was developed by him and some people from the business AI lab where they have it drawn during the class and the idea is that students can ask questions, but everyone will see the questions that is being asked and is use it as a helper that is like an order that is trying to answer questions live. Or also say someone asks a question to the lecturer and they don't know the answer, the lecturer can actually pass it back to the chat.

Kriti Khullar:

Thanks Andres, we'll stay with you for the next question. It's a two-part question. Is LLM an essential part of the education program and will ChatGPT and its applications be added to the curriculum?

Andres Villegas:

Okay, so you're asking about LMM. So, let me for the audience clarify what LLM is.

So, when we say LLM, that's Large Language Models and that's the technology that's behind this chatbot. So, when you say Large Language Models, it's the technology that is behind ChatGPT or any other text-based generative AI. So, this is an interesting question because you can see this as being two sides.

One is should we learn about the technology per se, how it's developed? And that can be contentious whether we should learn it or not. And it might be very specific to depending on the profession. And then the other one is how to use it. And this is going to be something that as we've seen in only nine months that it's going to become prevalent anywhere in life. So it's definitely going to be in education.

Now, whether we do it explicitly like this is a class on this, it's another thing. I don't expect that we're going to have that, but we're going to be using it and we're already using it quite a lot alongside other things. So you see in lectures, you might call the ChatGPT to help answer. You might ask students to use it as part of their answer. So, it's more like learning how to use it, but perhaps not directly, but as a helper to do other things the same way you use a calculator or people use a computer or use other things.

So, I would say it's going to be there, but I'm not that sure that it's going to be directly, this is a class on this topic that is LLM or ChatGPT or something like that.

Kriti Khullar:

Jon, what do you think?

Jon Shen:

Yeah, this is a fantastic question.

I think we as a society do need to have greater understanding and education about AI. And so not just LLMs, not just generative AI, but AI in the variety of forms it can take. And I've seen that there was some draft, I think national AI in schools framework that was being put forward a couple of months ago. And so we'll see how that goes and where that lands and what that means for education in schools and even primary schools and high schools. But the reality is because these tools are becoming so prevalent, we need people to have greater awareness of them.

And that means the risk and limitations of these tools, how they should be used in appropriate fashion, what can go wrong, including some of the inherent biases and problems that not so nice output that can be produced by these various tools.

So having this kind of information or this education being provided at younger ages I think is essential so that as people grow up, as they go through their teen years, they go into adulthood, they're at least aware of them and the challenges and problems that can arise from them before they suddenly get surprised by them in their later life.

Kriti Khullar:

Thanks Jon. Amanda, what are your thoughts?

Amanda Aitken:

I agree with Andres and Jon that gen AI tools are here to stay. And so I think educators everywhere will have a role or need to have a role in helping students learn how to use them and how they should use them in the workplace, including what can go wrong and how to avoid that.

On the side of whether you need to learn about the Large Language Models or whatever other technology used for other types of gen AI tools, I think it will obviously depend a lot on who it is that's learning, and even in the actuarial space. I guess it's a question almost for each practice committee of the Institute in Australia, will your members need to be involved in building these Large Language Models or do they more need to focus on how they'll actually use them?

Kriti Khullar:

So, Amanda, how can you tell if someone has used AI tools in their submission?

Amanda Aitken:

Yeah, I think that's another great question and I can see the lot of my students have maybe struggled with that a little bit over the last semester since some of these tools have been available.

I know a lot of them are using it in different or using these tools in different ways. So some of them are using it to help with their grammar. Some of them are using it to help generate ideas. Some of them are using it to help them code, and some of them are perhaps using it to just write the whole answer for them.

So yeah, ultimately, at least at the Actuaries Institute, a student needs to attest that everything they submit is their own work.

Kriti Khullar:

Thanks Amanda. Jon, your thoughts?

Jon Shen:

I agree with the points that Amanda has made.

It's so important for students to take the responses and really think about it and most likely rewrite it in their own words because first of all, sometimes the responses from ChatGPT can be a bit over the top or a bit verbose.

It's got some interesting linguistic qualities and it doesn't really reflect how we as normal people, as laypeople or even as technical practitioners might necessarily write text. So from a stylistic perspective, it's often useful just to take the text and rewrite it anyway, which then allows you to then think further as you're reshaping that text, is it correct? Is it doing the right things? And so on, which then goes to the point Amanda made about making sure it really is your own work legitimately.

And I think that's really where students will start to get additional benefits because if you assume for a moment, let's take a crazy world where every single student just takes the answer directly from ChatGPT, then everyone looks the same almost.

Whereas if you're someone that's taken that and used it to augment your own thinking, then most likely you're going to have a high quality answer, high quality response compared to the pure baseline. And that's going to serve you well both as a student in the education world, but also as you go into professional practice.

Amanda Aitken:

Jon, I really like that point you made about augmenting your thinking because I think that's another critical point about using these tools that if you are a learner and everyone is a learner, then what you want to do is just make sure that you read the output or view the output and use it to learn yourself rather than just using it as part of your say, assignment submission.

And the other point that I left off before is be transparent. So don't try to cover up that you've used gen AI in your answer, if that's what you have done. Let someone know that you haven't or source it or reference it appropriately or just make a comment that, "Actually, I use gen AI to help me think about this answer." And if you don't feel comfortable being transparent in that way, then you've probably used the tool inappropriately.

Kriti Khullar:

Andres, I'll throw onto yourself.

Andres Villegas:

Yeah. No, I think this is a great discussion, but first, if you're a student, remove that, "I am language model," when you paste your answer because you see some people submitting that.

But in general, I think a lot of the use is also the onus on us as educators to guide them, especially at these initial stages where everyone is navigating it, what is allowable and what is not?

So, in my case, I'm very open about the use. So what I ask my students is when they have an assignment, they always have to put an appendix where they say how they have used generative AI. And the purpose of that is for them to be transparent and also for me as an educator to learn how students might use it because there might be ways that they use it that I will never anticipated.

Kriti Khullar:

Jon, for the next question. With the emergence of advanced language models like ChatGPT, how do you envision the role of actuaries in the future?

Jon Shen:

My personal view is that data science actuaries need to understand how generative AI models work, how large language models are trained, and how to fine-tune these foundation models for domain or the organisation specific applications.

Now, some people might say that's a controversial view given how the technology is relatively new, but given there's so much industry focus and community focus on these tools right now, it's hard to believe that it'll be something new that rises up to take over it in the short term, the end term.

So I think it's going to be valuable for people to understand these tools, and there are two forms, the benefits take in understanding those tools.

Firstly, for data science actuaries, if you understand the technical details behind these models, then obviously you can build and configure them to best address the problems you're trying to solve from a day-to-day perspective when you're at work or in your different organisations.

But then going further than that, if you understand the tools and the risks and limitations behind those models, then you're going to be best placed to set guardrails and risk mitigation to ensure the tools are being used appropriately and the tools are doing the right thing in service of your organisations, your use case, your customers, and the broader community.

Now, more generally for actuaries in practice, again, it's going to depend if you need it in your day-to-day work. If you don't need to use the tool, then there's probably less value in you learning all the technical nuts and bolts. However, I do believe that there's value for actuaries to be at least aware that these tools exist along with their risk and limitations so they can have an informed conversation with others that might ask about these things in day-to-day work.

Kriti Khullar:

Andres, what do you think?

Andres Villegas:

No, I agree with Jon, especially in this part about putting guardrails about the use.

And in general, I think actuaries are very well-placed not only in generative AI, but in AI in general because we have certain skills that is not that common to see in the same professions. Having some technical understanding so we can see how they are built, how they work. Also, how to use them, and also what the implications. And actuaries have this understanding of the ethical, the fairness implications where I think is going to make us even more valuable because developing the tools is something someone can do it, but looking at the implications on how we implement, be it for business or for society, that's where the core skills are going to be more important.

That's something that we cannot outsource to the machines, but that we probably will have to do as actuaries or other professions, but actuaries playing a key role.

Kriti Khullar:

Amanda, over to you.

Amanda Aitken:

I think that gen AI tools are in some ways not dissimilar to many other tools that actuaries have had to work with and that have been introduced into the actuarial toolkit over time.

So going back however many years, we didn't even have calculators. We had to look up values in tables, then we've got calculators, then we had Excel, now we've got gen AI tools.

We're constantly being given more tools to help us in our work, and I think throughout the evolution of actuarial work, those tools have just helped us to free up space so we're not spending so much time doing more menial tasks and we have more time to evaluate the output of the models that we build, for example, so that we can more play the role of expert advisor than technical boffin. So, I see these tools as contributing to that same path of evolution for the profession.

Kriti Khullar:

Amanda, I'll stay with you for the last question for today's podcast. What skills should actuaries brush up on or build to remain relevant in this rapidly changing landscape?

Amanda Aitken:

I think that a lot of actuaries are probably a lot like me a few years ago where there are all these new terms being thrown around and they felt very strange and very far off.

And so probably the best thing that an actuary can do is just gain some basic understanding of different data science terms and tools. So, something like the Actuaries Institute's microcredential in the Data

Science Applications course could be a great place to start to just understand what are deep learning models that are part of the technology behind Large Language Models? What is natural language processing? What do we mean by supervised versus unsupervised learning?

So, a lot of those concepts, once you've covered them, for example, through the microcredential that we offer, I think will give you enough of a base understanding to then build towards understanding some of these more modern tools like technologies that are used for gen AI.

Kriti Khullar:

Before we wrap up, Jon, Andres, is there anything you would like to add?

Jon Shen:

A few months ago, maybe six months ago, there was a meme floating around that prompt engineering is the new profession to be in given all the hype around ChatGPT and large language models.

And while I think that hype has died down a bit, and it probably was more of a meme, the point stands that knowing how to prompt a large language model and iterate over a prompt to get a meaningful output is still going to be a useful skill for us as practitioners or even as laypeople because no matter what question you're asking, you're going to give it the first prompt. You ask it a question, it will return something and you might like the response and that's fantastic. Or you might say, "That's not good enough. I need to tweak the text or the question I've thrown at it in the first place." And being able to iterate through that process. It's just like something that we all know, the actuarial control cycle, taking input, getting output. If it's not quite working, then tweak it, try again.

So, the short summary of that is everyone, if you haven't already done so, should get an account, try talking to ChatGPT or another Large Language Model just to see what the hype is about and understand how these tools might make your life easier even in everyday tasks.

Kriti Khullar:

Thanks everyone. Listeners, we hope you enjoyed this discussion with Amanda Aitken, Andres Villegas and Jon Shen.

If you're hungry for more conversations and insights, we've got you covered. There is a treasure trove of Actuaries Institute podcasts waiting for you at our podcast hub, which you can access via the podcast section of Actuaries Digital, or via all mainstream podcast apps.

If you'd like to continue the conversation with us on social media, join us on Facebook and LinkedIn @ActuariesInstitute, or you can reach us via Instagram @actuariesinst. Please write in with your questions and comments on the show. We'd love to hear from you. I'm Kriti Khullar, bye for now.