



GenAI in Insurance

Ed Chanda:

Hello, everyone. Thank you for joining us today. Generative AI is a hot topic in the insurance space, and we're all interested to hear more about it. Kelly, why don't you just start with, what is generative AI?

Kelly Combs:

I like to start by thinking about artificial intelligence as a broader set of capabilities that are trying to mimic human behavior and under artificial intelligence, many are familiar with machine learning techniques and deep learning. And, now we have this field called generative AI, that's a set of unsupervised and supervised algorithms or models with language capabilities wrapped around it. So, what we're doing is we're taking some techniques that already exist and packaging them in a different way, that's ultimately bringing data science closer to business stakeholders in a way that they can more intuitively interact with and receive insights from AI.

Ed Chanda:

We've been talking about AI for many, many years – old AI, machine learning, those kinds of things. And it seems like, in the last 5 or 6 months, we've come faster than we have in the last 5 or 6 years. What do you attribute that to?

Kelly Combs:

I almost think, there's this excitement, and we're definitely in the hype curve for generative AI, no doubt. But there's excitement in the fact, and maybe more of an appetite to adopt generative AI or to use generative AI, because we've already gone through the low code or automation challenges, some organizations as early as 7 years ago or 8 years ago. There's a familiarity with using tools that augment how we work as humans and business decision making, that's making adopting gen AI faster and accelerated. And, as I mentioned before, as we move data science closer to business stakeholders, it's not that we're eliminating the data science element or that you need to be a deep technologist and understanding how the

technology works, but we are packaging algorithms and AI in a way that's more intuitive for normal users to use and interact with. And, the capabilities in itself are fairly broad of what generative AI can actually do. That there's this idea, just like low code, we can use it across multiple functions, we can use it to augment the individual and how they work, and there's some transformative use cases at the enterprise level that we could do at relatively low cost to start. The caveat being, how these services are deployed by vendors and how their package is probably going to change a lot in the next 6 months to a year. But, it's a relatively low barrier to entry to start. I know we'll probably get into some of the risk side of it, but I think that's helping accelerate adoption in this area. And really, we're seeing across clients, and industries that everyone's excited about it. It's not particularly one industry.

Ed Chanda:

Sure, and we get more attention from a lot more people upfront than we usually do. Although, we do tend to overestimate how fast things are coming. I know we're very excited about it, and the old saying that humans tend to over overestimate what's going to happen in two years and underestimate what's going to happen in 10. I wonder if we've solved the data problem at the bottom of all of this. You said we've been working on this for some time, but I still think there's a lot of data work to do. Will AI be slowed down at some point while we try to get it a better source?

Kelly Combs:

Yes. AI, and generative AI in particular, is the newest, shiniest tool in the toolbox for business stakeholders. We still have challenges with traditional AI, again, what organizations are building internally or using algorithms for, that journey they've been on for a couple of years as well, in parallel to automation. And what we're saying is, hey, the tools can do what we want them to do. But is our data structured in a way to deliver the right insights? So, in parallel, now we want to add Gen AI as

one of those tools that can augment decision making. We're now also seeing a lot of organizations refocus on, is our data conducive to using these tools in the right way to deliver value? How fast and quickly can we adopt these tools? At what cost is it going to take? So, there's definitely a renewed focus on the data element. And, a lot of questions on what's the right framework to use for build versus buy. Again, automation was pretty clear in terms of licensing tools and some of the big vendors out there. There's ways you can build AI and use hyper scale or workbenches to deploy and develop your own models, or use pre-built models or open source tools and libraries. Generative AI is kind of exploding AI as a service concept and probably challenging the framework of which use cases make sense to use which tool. There's interest in using Gen AI probably for everything to start, but then we need to step back and say strategically, what are the patterns or capabilities we're trying to achieve based on the business problem? Is Gen AI the right tool or is it one of these other options that we may already have?

Ed Chanda:

What are some examples of what generative AI does well? Where are we seeing it go first?

Kelly Combs:

It's really good at things like creating content and generating new content, creating synthetic data, writing code, visualizing information. It's really good at explaining, editing, comparing type tasks. If you think of researching, querying or trying to benchmark data, whether it's summarizing, synthesizing or transforming something into a narrative form. It's also really good at acting as a coach, helping provide options, look at different scenarios, predict outcomes, and then provide a preliminary recommendation to a user that then can interpret and make the ultimate decision. So, the basic, create, summarize, explain, edit, compare tasks, generative AI is really good at. And, it's trained on public information. Chat GPT 3 was trained on data through 2021, Chat GPT 4 has more recent data, and it is across the board and public sources. The generalized applicability of using it is good, that being said, if you want something that's very domain specific, or very tailored to your use case or providing insights very specific to your organization, there's ways that you can connect generative AI to your internal corpus of data and ultimately train it over time to learn more about what your organization does. But again, there's different ways to deploy it and different trade offs for using it in a more generalized sense versus a more narrow, niche or domain specific context, which is going to take a little more work to get the accuracy right.

Ed Chanda:

We're still at this stage talking about a technology enabled human, not the human giving up responsibility. That's correct?

Kelly Combs:

That's right. The way we're seeing clients deploy this is two different scenarios. One is, how can we use generative AI for the individual? Just like again, low code and automation. How do we augment the way that our employees work? How do we use it to make us more efficient? To do those stare and compare type tasks or research tasks that take a lot of time, just to boost outcomes or get to insights faster and increase productivity. The second is the more enterprise use cases that are more transformative, which is maybe requiring us to train on our internal data to provide more accurate or more specific insights, but still a human in the loop in both cases. So one is, how can I better create a PowerPoint presentation on financial forecasting, and how can I just gather the information faster? Whereas the second, how do we improve financial forecasting across the finance group? There's different implications around the data and how you deploy it, but both require humans to say, is what I'm receiving in the context of what I'm trying to do accurate? There are limitations on the output, so we definitely emphasize augmenting humans versus replacing or deciding for humans.

Ed Chanda:

That's very consistent with what I'm seeing in the insurance sector specifically. We're in the business of taking on risk and making judgments, and I could see a role for AI someday in the underwriting process, but there's a lot of risk to getting that wrong. Where we're seeing it go first is to marketing and communications, being able to get more consistent, crisp and clear communication going with our customers. Or, looking at a compliance function where it's bringing back what it finds to a human, and the human decides whether or not to take action. Nothing as bold yet as underwriting where we're binding the organization to a contract based on what AI thinks. Maybe that's next?

Kelly Combs:

We're seeing a lot in marketing and sales, and how we personalize offers, that point of interaction, promotions, or even just communications. How do we make them more accurate and tailored to the end user? Which again, the generation and content construction component of generative AI is really good. But, like you said in certain areas, you might say, we've got a proprietary model that we use to look at risk or to assess how we want to approve claims or certain areas, and the models we can build internally are more effective and precise versus where we may use Gen AI. Another area I'll speak to is in call centers, or even looking at claims management. How can we augment what an agent is doing to make them more efficient? It's not making the decision on the claim or making the resolution through an interaction with a customer, but it's helping surface up related information

about that claim [such as] potentially recommended next best action to work through with the client that you're talking to or the consumer with an inquiry. What we found in that scenario is the productivity gains for early employees within an organization is huge. If you're an experienced agent, for example, the net impact of Gen AI is neutral – it doesn't really boost performance per se because you've got that domain context of working in that area. You've seen certain types of claims, you know how to process them. But for early joiners within an organization, it is helping bridge the training gap and bringing the productivity of the group as a whole up. So, how do we augment and improve efficiency in claims management call centers/customer interactions? That's another good use case that we've seen a lot of organizations start to experiment with.

Ed Chanda:

That makes a lot of sense, but I do worry that, are we breaking the apprenticeship chain? Are they depending on AI, or are they learning from AI? So, after 20 years in the business, I know how to do it. If I can get to that level in a year and a half because I'm using AI to help me, am I in a position to know that AI is doing it correctly or that's doing the right thing? How do I learn?

Kelly Combs:

It's an interesting trade off between, what do we learn through experience and on the job training, and how structured is that? How much can we use technology to augment the standardization of decision making? In the field in general, is this idea of the prompt engineer, and how do we prompt generative AI in a way to get the same and consistent output? I can envision a world where you have agents prompting generative AI, asking similar questions, but may get different results back. You still want to standardize how recommendations are being made, how the process is being facilitated. A lot of that today is through job aids, policies, or workflow tools. But, if it's generated by Gen AI, we need to think about putting applications in front of generative AI to make sure people are asking the question in the same way to get a consistent output back. So that's something else I think about as well, in terms of the skill gap and experience, and how do we make sure people are learning the right thing in the same way.

Ed Chanda:

You mentioned call centers, I wanted to come back to that. Are we almost to the point where call centers are going to be a pleasant experience?

Kelly Combs:

It may be too soon to tell. I don't know, but like I said, we are seeing big savings there. Just from internally within the organization, 40% to 60% are expecting productivity

gains in customer service. It's a huge opportunity to upscale or up level productivity across the board. How that translates into, is the human interaction and experience component faster? Or, can we get to root cause analysis for your problem faster? I'd like to think the answer is yes, but I assume it's still probably not going to be entirely frictionless.

Ed Chanda:

Let's talk a little bit about where the vulnerabilities might be in this new model. Where can things go wrong and what new risks am I introducing that, maybe I haven't had to deal with using previous technologies?

Kelly Combs:

I like to think of it as, every organization should probably be thinking about responsible AI or AI governance more broadly. Whether they're building a lot of AI models internally, using third party services, or using a chat agent assistance, I think it's a good practice given where AI adoption is at more broadly, to have policies and procedures. Many organizations do [have] principles on what good practices are for AI, and have an ethical stance on, what permissible data do we want to use for the purposes of AI? Where do we want to use certain algorithm techniques? Where do we not based on how transparent they are, how complex they are? My first advice is always, baseline, there should be AI governance established through policies and procedures. The nuance around generative AI is, it's almost like a double click into that policy or procedure that may exist, really focused on four key areas. One is security. Security is a growing concern. But if I think of AI as a service, we still have challenges around potential manipulation of inputs or outcomes. Can we do data injection or data poisoning and try to uncover information we normally wouldn't be privy to? There's a lot of fraud and counterfeit happening. Again, even trying to manipulate Gen AI and things like adversarial attacks. So this idea of security and how AI serve to endpoints, and how we secure those endpoints, and how we secure the data sharing element is going to be of increased focus. Of course, contracting so generative AI is packaged as a service. Not many organizations are building Gen AI from the ground up internally, they are leveraging it from a third party vendor. Thinking through, who owns intellectual property and copyright, which is still being sorted out legally. So, if I receive generated content and then I adjust it, do I own the IP around that? Is it public information? Is it the vendor that produced the initial results? That is still being sorted. But other areas, rights to audit, data rights, and understanding the roadmap of these vendors. The contracting in itself is a huge space we could probably dig into. But like I said, that's going to change a lot. The two other areas, I'll hit on that. I mentioned the outcomes, so are there limitations on the accuracy? We don't know how the models are trained. We don't

know what data is being used. How do we know what we're receiving is accurate? And then data kind of being the fourth one around data privacy. Can we turn off how we log data on the back end? How do we make sure we're not sharing sensitive data? Those are the four big areas of focus for Gen AI, that I would say, if we aren't thinking about and you're experimenting, those are the areas to double down on from a governance perspective.

Ed Chanda:

You mentioned data poisoning, changing what's going into the model, it all keeps coming back to the same question – how do I tell how the answer was arrived at? How confident I should be in that answer? And then, as we try to use it in certain spaces, I mentioned underwriting before, but even if we try to use it in audit, we need to be able to prove that we have some basis for understanding the answer that we're coming to, not just accept the answer. I think that question is going to keep coming up over and over again as we try to introduce it into new use cases.

Kelly Combs:

Or even a scenario I've seen, you can prompt and ask for a recommendation, and then if you have the knowledge to say, well that doesn't quite seem right, and you ask the question in a different way, generative AI could come back with a different answer. How do you interpret and

effectively challenge what generative AI is giving you and not taking it for surface value? It's hard to say who's more knowledgeable – is it the human worker with 20 years experience to make the ultimate call or is it generative AI, which is trained on millions or trillions of attributes of data? That's going to be where organizations have to focus thresholds, controls, processes around. How do we get comfortable with what we're seeing, and ultimately, what is the right and most accurate answer in the context of how we're trying to use the information?

Ed Chanda:

I know that we could probably peel back the next layer and go another hour here, but I think this is where we're going to end. Thank you so much for sharing your thoughts with us and we look forward to talking again in the future.

Kelly Combs:

Thank you for having me.

kpmg.com/socialmedia



Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

The information contained herein is of a general nature and based on authorities that are subject to change. Applicability of the information to specific situations should be determined through consultation with your tax adviser.

The following information is not intended to be "written advice concerning one or more Federal tax matters" subject to the requirements of section 10.37(a)(2) of Treasury Department Circular 230.

KPMG LLP does not provide legal services.

© 2023 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization. USCS005399-2A