

Consumer Finance Monitor (Season 6, Episode 16): A Close Look at Generative Artificial Intelligence and What it Means for the Consumer Finance Industry, with Special Guest, Alex Johnson, Founder and Author of Fintech Takes Newsletter

Speakers: Alan Kaplinsky and Alex Johnson

Alan Kaplinsky:

Welcome to our award-winning Consumer Finance Monitor podcast, where we explore important new developments in the world of consumer financial services and what they mean for your business, your customers, and the industry. This is a weekly show brought to you by the Consumer Financial Services Group at the Ballard Spahr Law Firm. I'm your host, Alan Kaplinsky, the former practice group leader for 25 years, and now Senior Counsel of the Consumer Financial Services Group at Ballard Spahr. And I'm very pleased to be moderating today's program. For those of you who want even more information, don't forget about our blog, consumerfinancemonitor.com. Goes by the same name as our podcast show. We've hosted our blog since 2011 at the very same time that the Consumer Financial Protection Bureau became operational. There's a lot of relevant industry content there.

We also regularly host webinars and subjects of interest to those in the industry. So to subscribe to our blog or get on the list for our webinars, you should visit ballardspahr.com. If you like our podcast show today, please let us know about it. Leave us a review on Apple Podcasts, Google, or wherever you obtain your podcast. Also please let us know if you have any idea for other topics that we ought to explore on our podcast show. We're always looking for fresh ideas. And also if you have any speakers for our show or guests that we should consider inviting to be on our show. Let me tell you very briefly what we're going to talk about today and then introduce our guest. The topic is artificial intelligence, but more specifically generative artificial intelligence.

Artificial intelligence is already in use in the consumer finance industry. It's often used in connection with underwriting. There are a number of companies that specialize in creating artificial intelligence software that is in use by a lot of credit unions, banks, and others that are in the consumer finance industry. It's a very controversial area, I would say. The Consumer Financial Protection Bureau, CFPB has not been particularly enamored with the use of AI because it believes that it has fair lending problems associated with it. And the CFPB has issued a number of statements and has issued in various forms on their blog and speeches where they warn people that are using AI for underwriting, that that doesn't exempt them from any of the consumer finance laws that might otherwise apply.

Like the Equal Credit Opportunity Act says you can't discriminate against a protected class, that if you turn somebody down for a loan, you've got to give a reason why they've been turned down. And that has often been a very challenging task for companies that are using AI. Well, we're going beyond AI today. We're going to the next, I guess I could say the next generation of artificial intelligence. Those of you, most of you probably have heard about ChatGPT. That's a type of generative artificial intelligence. You may have heard about it from your kids who have wondered whether they instead of having to write their next term paper, they can just tell their computer that they want a 10-page paper written about something, whatever the topic may be. And within minutes a term paper will be produced off the printer, and that will be that.

Well, that is one potential use cases. But today we're going to talk about potential use cases in the consumer financial services world, and I can't think of anybody more qualified to talk about this subject than my guest, Alex Johnson. Alex is the founder and author of a publication called Fintech Takes. It's a media publication and a podcast that's focused on analyzing the intersection of banking and technology, really FinTech. He has more than 15 years of experience in the financial services industry, including stops at Mercator Advisory Group, Cornerstone Advisors, and FICO. FICO of course is, those of you that

are in the industry know that that is the company that produces a credit score or a FICO score. So first of all, Alex, very warm welcome to you. A pleasure to have you on our show today.

Alex Johnson:

Thank you so much for having me. It's my pleasure.

Alan Kaplinsky:

Great. Alex, got a number of questions for you, because this subject really is fascinating to me. This technology, which is called generative AI. Can you define what that is for us and how is it different from what came before, which I guess was the ordinary AI that I talked about that some companies and banks are using right now to underwrite loans?

Alex Johnson:

Absolutely. So as you noted, AI and machine learning are not new, either as a general field of computer science or specifically being used in financial services. We've used versions of AI and machine learning in lots of different contexts in banking for a long time. Generative AI, which is this new and emerging category of artificial intelligence techniques, to really simply define it, it's basically machine learning techniques that are designed to generate, so hence the name generative AI, original content based on a given set of input data. And the content that gets outputted from these generative AI models can be text-based. So like you mentioned, ChatGPT is a chatbot built by OpenAI. It uses generative AI to basically understand text prompts that it gets from users and then respond in a interactive, fun and useful chat format, giving you back text answers.

And those can be relatively short answers to specific questions, or as you noted, it can write term papers for students in school, based on these prompts that it's given. There are also versions of this same type of generative AI technology that are being used to create images. So folks who are listening to this might be familiar with a tool called DALL-E, which you can use to basically, again, with a text prompt, generate images of pretty much whatever you wanted to see. So if you wanted to see a grizzly bear wearing the top hat attending a musical in the style of an impressionist painting, you can basically type that into this tool and it will generate pixel by pixel, that exact image that you asked for.

Alan Kaplinsky:

Great for creating emojis, huh?

Alex Johnson:

It is. It absolutely is. There's just so many cool potential outcomes. As you might imagine it also is freaking out certain people. For example, on the last one I just talked about, if you're an illustrator or someone who creates graphics for the internet, this ability for anyone to just type something in and generate a graphic is a little scary maybe to your livelihood. It's a very new and exciting field. To answer your question about what's different about it and how it works. The big breakthrough here that sits behind these generative AI techniques, are something called large language models. And basically what those are is these models that have learned to predict the probability of the next word in a desired output response based on the context of all of the preceding words.

And so it's basically this thing that gets really, really good at predicting what it should say based on the prompts that you gave it and the words that it has already typed. And so if you've used like ChatGPT, you'll notice that as it's responding, it's responding in real time and it looks like it's typing out a response. It is basically predictively going through and figuring out what is the right response based on the prompt that I was given. And the way that it works, and it's really very interesting, is these models are all trained on these huge, huge amounts of unstructured data. And that's really the thing that makes these generative AI techniques different than what's came before, is that they're very efficient at churning through a lot of data. So in the case of ChatGPT, the folks who built it actually basically just went out and scraped almost all of the words on the entire internet.

So just everything that we've been writing on the internet since it's existed, they had this huge corpus of texts that they pulled together about 300 billion words or so, and they basically had this model just churn through all of those words and learn all of

the relationships between those words. And the relationships between those words and its ability to understand the context between all of those different relationships, that's what allows it to output these startlingly human sounding responses based on these prompts that you give it. And once they figured out how to do this from a unsupervised learning perspective where it just lets the model just run through all these words and figure everything out, then they started to figure out, well, if we have this technique, now we can apply it to images, we can apply it to music, we can apply it to 3D modeling. It's the same basic idea of generating some type of output, some type of content based on the input that you give it.

Alan Kaplinsky:

So you mentioned music. Is it actually able to create new music?

Alex Johnson:

It is, yeah. Basically same thing as with the words and with the images, you can feed this model a corpus of existing music and audio files. Again, just a collection of everything that human beings have created. And it can teach itself the relationships between notes and between different harmonic patterns within songs. And then you can generate new songs, entirely new original songs based on its understanding of the relationships between different elements of existing songs. And one thing that's particularly important to understand is, these models get built in an unsupervised fashion based on a huge set of training data. But then what you do is you then go in and do some refinement of that model. So you use supervised fine-tuning of the model where you give it a smaller data set that's more specific to the use case that you're trying to do.

And then you might even have a human being at the very end who provides reinforcement learning on top of that to just fine tune the model a little bit more, because a lot of the things we're asking these new tools to do, respond like it's a human being typing on the other end of this thing. Write a pop song in the style of Korean pop artist. Generate this image that looks like a five-year-old drew it with a pencil. That requires a lot of nuanced understanding, and these models are pretty good, but you need that last level of supervised training to help teach them the last few nuances of how to do whatever it is we're asking it to do.

Alan Kaplinsky:

Wow. Well, anyway, let's talk about, let's bring it back to the more mundane areas where it might be used such as in consumer finance. I want to run through a bunch of specific use cases, and I'd like to hear what you think about the potential of using generative AI in these areas. Okay? So we'll start with marketing.

Alex Johnson:

On all of these use cases, one thing that's important to keep in mind is that you want to separate out what these generative AI tools are good at and what they're potentially not good at. And so one I think very simple framework for breaking apart all of these use cases is to ask what the stakes are for the use case. So how critical is it for this particular use case that we not make a mistake? And the reason that that's important is that these generative AI models are predictive. They try to guess what the output is that you're looking for, but they're not perfect because anything predictive, it's going to be in the ballpark of being accurate, but it's not going to be accurate every single time, 100% of the time. And if you've played around with any of these tools, you'll know they screw things up.

For example, with those tools I mentioned that will generate images based on prompts. When you ask it to generate an image of a human being, strangely one of the things that's not very good at is drawing hands. And so you can look at the hands of these pictures of humans that this tool builds and it'll have too many fingers or the hand will be distorted in some way. So they make mistakes. And the reason that's important is that obviously as your audience knows very well, in financial services, there are definitely a set of tasks within any family of use cases where the stakes are just too high to make a mistake. And so I think when you talk about marketing, you can separate marketing into high stakes and low stakes tasks.

And I think a good example of a low stakes task where something like ChatGPT or a generative AI tool could be really useful would be today we try to send personalized emails and marketing messages to consumers, and there's a human being on the other side of those emails writing each one of them. And obviously there's a limit to how many emails a human being can

write. So you have to have these very broad segments where you go, okay, this is a segment of our customers, we're going to write an email that's tuned to them, and then we'll have another segment and maybe you do a dozen segments or a couple dozen at most. But that's a lot of work for a human being to do. With generative AI you could personalize marketing emails to consumers almost down to an individual basis, because the cost of producing those emails goes down dramatically when it's not a human being writing every one of them.

So that's I think a very good example of a low stakes task that you would be pretty comfortable turning over to a well-trained generative AI model. On the other end of the spectrum, within marketing, there are obviously things that we do that are much higher stakes. Like when we get into things around pre-screen or pre-qualifying customers or anything that starts to merge into underwriting or credit, you really want to have a high level of certainty that you're doing that right, that you're complying with all the applicable regulations and that there's not some generative AI tool that's like guessing rather than knowing what it should be doing there. I think within marketing, the things that are more about personalized messages or some of those very easy chatbots where they're just answering very simple questions for a customer, those can be a good example of where this can provide value.

Alan Kaplinsky:

Okay, that's interesting. Well, as you pointed out, I can't think of an industry as highly regulated as consumer finance at the federal level, the state level, the laws differ from state to state. We'll stay with marketing for a minute and I want to move on to some other use cases. But do you foresee a future where generative AI will be not just predictive, but it will be perfect or at least as perfect as a human being can be?

Alex Johnson:

Yeah, that's a good question.

Alan Kaplinsky:

Because human beings make mistakes too.

Alex Johnson:

They absolutely do. They absolutely do. I guess two things on that. One is that the generative AI by its nature will always be probabilistic. And so if you think about the distinction between a slot machine and a calculator, generative AI, just because of the way it works and the way it's designed and the way that it essentially is just a really good guesser, its guesses might get better, but it's always going to be guessing. And so I think it can definitely improve the way that we train these models can get better. The amount of unstructured data that we feed into them to train them can get bigger. And all of the companies that are working on generative AI, working on making these models better and more accurate. But it's always going to be probabilistic in nature, and so they're never going to be 100% accurate.

Now, to your point, human beings are not 100% accurate either. One thing though that I've observed across a lot of different areas where we're starting to apply more automation and more AI is, there's a huge gap between what we as a society are comfortable with human beings screwing up and what we're comfortable with machines screwing up. And you see this with self-driving cars as an example. Self-driving cars, the technology is actually on average at least for certain instances or use cases safer already than it is for human beings. The problem though is that even though it's a little safer than human beings, it's not so much safer that we are comfortable turning ourselves over to that technology and putting our lives in its hands.

Financial services is a little less life or death than driving down a highway, but as you pointed out, it's similarly high stakes. And I think even when these models get maybe on average better at guessing than human beings are, and they might already be there, it's going to be a while longer before we're comfortable actually handing the reins over to it.

Alan Kaplinsky:

Yeah. All right. We talked about marketing. What about a big area for banks and non-banks as well, fraud management?

Alex Johnson:

Fraud is a really interesting one because that is pretty high stakes for the most part. And so if you have models that you're using to predict fraud or you have rules that you've put in place to screen out fraud or deal with fraudsters, you probably are pretty comfortable with the level of predictive power that those models have. And more importantly you're very comfortable being able to explain how those models work. And generative AI, generally speaking will make it more difficult to explain why the model works the way that it does. And so I think in the case of fraud, when it comes to actually trying to predict fraud or trying to identify fraud, particularly in very high stakes situations like account opening or when a customer swipes their card and you have a couple milliseconds to figure out if it's a fraudster or if it's a legitimate customer transaction, those use cases I think are going to be a little too high stakes potentially for generative AI, at least in the short term.

However, I think an area that financial services companies need to think a lot about with generative AI and fraud, is that fraudsters are going to start using it. So one of the things you see in fraud is whenever we invent some new tool that companies start to get excited about, fraudsters also get excited about it. And a lot of times they end up beating companies to the punch and start using these tools to commit fraud and to be more effective at stealing money than the companies are in learning to adopt it for legitimate use cases. And so to give you an example, phishing and other types of social engineering scams are rampant inside of financial services right now. You get a lot of fraudsters that target consumers with Zelle accounts and they'll use social engineering to essentially convince the consumer to authorize a transaction or a payment to the fraudster that then can't be recovered and that maybe the customer ends up being liable for.

Those types of social engineering scams are only as effective as they are because you're convincing the customer that the request you're making is legitimate. And if you've ever gotten a phishing email or something, if you know what you're looking for, you can tell they're fake, right? Because they're spelling errors, they're weird syntax things. It just doesn't sound like a human being wrote it.

Alan Kaplinsky:

They're getting better though. I get these things almost every day. It's unbelievable. And it purports to come from a bank. Sometimes it's a bank I do business with. Sometimes it isn't. Sometimes it purports to come from Amazon. We're shutting down your account. They can look scary.

Alex Johnson:

Absolutely. And so the better those get, the better they work in terms of convincing customers to authorize things that they shouldn't authorize. Generative AI could be a game changer in this respect because it would make it much easier for fraudsters to be really compelling. It could write an email with no spelling errors and no easy indicators that, hey, maybe this is suspicious. It could even go so far as to we're starting to see generative AI tools that can dynamically generate voices, sounds. Right? And so you could get a call and instead of it being a fraudster that maybe doesn't speak perfect English or that there's some cue on the other end of the line that maybe this isn't legitimate, it could be a really great human sounding call that responds perfectly, that's a machine talking to you on the other end that's been employed by a fraudster to convince you it's a bank account rep from your bank calling you.

And so the attack profiles from fraud I think are going to get a lot more sophisticated because of generative AI. And the way I would look at it from a bank's perspective is, what do we do to identify when we're being attacked by generative AI or when our customers are being attacked, because I think fraudsters are going to get to this stuff pretty quickly.

Alan Kaplinsky:

Wow. Wow. Okay. Credit underwriting.

Alex Johnson:

Credit underwriting, as you noted, is probably going to be one of the last frontiers where we see this being used, at least in terms of turning the decision over to it. I do see some scenarios where generative AI could be used to assist human underwriters. So as an example, if you are underwriting a large jumbo mortgage or something that falls outside of the

traditional boundaries that you're allowed to approve as a bank, the process might be that the loan officer needs to write a memo outlining the situation, the request from the customer, the risks associated with it, and then present that memo to some type of board or committee that approves those loans that are outside the normal parameters. Putting that memo together is very time-consuming, and you could use generative AI to do a first draft of that memo to make that underwriter or loan officer's life a little bit easier.

So there are things like that where generative AI can be a co-pilot to an underwriter to help make them more efficient. But I don't see a world in which we let generative AI make underwriting decisions because again, it's not going to be quite good enough to guess in order for us to feel comfortable with that.

Alan Kaplinsky:

Sure. We're going through the progressive stages from marketing and credit underwriting. What about collections?

Alex Johnson:

Collections is another one that I worry about a little bit honestly, because obviously there are some actors in the collection space that are a little unsavory in their techniques to try to recover money, particularly the third party collection space. And I do think that you will see generative AI make its way into that space, because again, like the fraudster example that we talked about, if you want to try to convince consumers to do something, generative AI it's very compelling and it's very convincing, and it does a good job of generating outputs that drive a reaction or a behavior because it's predicting what's going to work best. I could see it being used in collections to reach out to customers to do collections treatments, to try to get customers to pay.

I expect that we will probably see some evolving regulatory opinions and maybe some rulemaking around that, because collections as your audience probably knows, has been stuck in the past in terms of regulation and rulemaking and what you're allowed to do. And can you do something as simple as text a customer about their debt, that kind of thing. I think those rules are going to have to continue to evolve to account for generative AI where it's not even a human being or a rule-based system that's reaching out to the customer. It's one of these generative AI tools that's talking to the customer dynamically.

Alan Kaplinsky:

What about customer service? Does that seem like an area where maybe there might be a better use case?

Alex Johnson:

For sure. Yeah. I think so. I think that, again, you can divide customer service use cases into high stakes and low stakes. I think for low stakes customer service requests, we can definitely use generative AI to make that better. A lot of banks already have chatbots that they use for basic customer service requests in order to self-serve those. I think we can use generative AI to make those chatbots much smarter and more engaging than they are today. I think that one is a slam dunk. I think we will see that fairly soon. There are higher stakes customer service interactions where you might not be totally comfortable with generative AI taking over on day one. So if are, say for example, if you're doing financial advising where you have a fiduciary duty to the customer to make good recommendations and you're taking on some liability for making financial recommendations, I probably wouldn't want to turn that over to generative AI day one.

But you could use it as a co-pilot to help assess, hey, what do you think maybe the best options for this customer would be? Which products maybe should we recommend? And then just filter that through a human being who makes the final recommendations to the customer. I also think that you could potentially see generative AI, and I wrote about this in my newsletter. I think you could potentially see it become a financial therapist in a way, because a lot of the unmet customer needs in the customer service space relate more to helping consumers understand their relationship to money and process their emotions around money. And most bank tellers or customer service agents aren't really trained to do that kind of thing. Even financial advisors are more about the dollars than they are about the emotions behind them.

So I do think that there is an element of financial coaching, financial therapy that these generative AI tools could potentially be really good at. So that's an area where I could see some innovation as well.

Alan Kaplinsky:

You mentioned financial advising I guess already, and I get the sense that that's a situation where it can be used as a co-pilot, right? Not to replace human beings.

Alex Johnson:

That would be my sense as well. And again, it's like back in the old days when they first invented elevators, no one wanted to use them because they didn't seem safe, and so you'd have an elevator operator who went up and down with the customer. It's not like elevators really needed that elevator operator to do anything special. They were more just there to make the person comfortable. I think financial advising might end up looking something like that, where the generative AI tool might be doing a lot of the work behind the scenes to process what the customer's needs are, to try to figure out initial product recommendations and maybe even to script out what the financial advisor could say to help engage that customer, but we're going to want to have that elevator operator in the elevator with them is my guess.

Alan Kaplinsky:

What about back office use where it does not involve direct interaction with the consumer? Would that be an area that might be a better use case?

Alex Johnson:

I absolutely think so. I think there's a ton of back office examples where you're just having human beings spending their time on very low value tasks that are not high stakes, they're not particularly high value. And candidly it's expensive to employ human beings to do all of that work, and it's mind-numbing work also. AI and in particular generative AI could be really good at some of those tasks. So if you have examples where you have human beings reading through big long documents and then summarizing them, or as an example, one thing that's mind-numbing would be, hey, review all of our contracts with technology vendors and summarize those contracts and what the main takeaways are from them so we can have a new strategy for negotiating with our vendors.

You could have ChatGPT or a generative AI tool go through all of those contracts, summarize all of the key points, and then give that to human being to do the high value work. But all of that back office drudgery gets automated away. I think that's the thing that you'll see much more than the stuff that touches the end customer.

Alan Kaplinsky:

Right, right. I guess it's fair to say we're in the very early stages of this technology. What remains to what needs to happen for these generative AI tools to be ready for primetime in the financial services area?

Alex Johnson:

The big thing is that they need to be fine-tuned for financial services. I mentioned before that the process of building one of these models is you start with a very, very large amount of unorganized text or input that you feed into it. And so it's very unsupervised. The data is unstructured and the model just churn through it on its own and teaches itself the relationships between the data. That's where all of these start. However, all of the things that actually end up getting touched by customers or getting touched by employees at companies have to go through a lot more of that refinement after that happens. And so I think that in financial services, because all of this technology is still so very new, we haven't done much of that refinement work yet.

And so if you're starting to think about what are good use cases for generative AI in financial services, what you're going to want to think about is, how can we take one of these large, somewhat unstructured models that's already been built and how can we tune it to our particular use case? And I think in financial services, that fine-tuning process will be significantly more intensive than it's been in other industries because we have such tight requirements around making sure this works, making sure we have a high level of confidence around it. So I think that fine-tuning process will take longer, it'll be more expensive, they'll be more resources put into it. But I think that's the next step for generative AI and financial services.

We see these generic consumer tools like ChatGPT or folks might have seen Bing, the search engine from Microsoft is getting a chat interface added to it. That's the same type of thing. Those general consumer use cases are a little easier to fine tune for because it's just more generic. But in financial services we have a lot more specialized work that needs to happen.

Alan Kaplinsky:

Sure. I've seen that a number of financial services companies have banned the use of generative AI by their employees. The other companies, I think Chime being a good example, plan to build their own tools rather than build on top of the ones that have been created by OpenAI or Google. What are these companies concerned about?

Alex Johnson:

Good question. I think there are two primary concerns. One, in the case of just banning these tools inside of these companies. I think JPMorgan Chase is an example of a company that banned the use of these tools within its walls. That's I think more about just in the same way that Chase or these very large banks will restrict their employee's use of different internal communication tools. They want everything to flow through a controlled set of tools that employees use. I think they have similar concerns around data privacy, making sure that they're not sharing any proprietary information or asking any proprietary questions, and then just all the different laws that they're required to comply with around how their employees communicate with each other, with customers or whatever.

I think that's one set of concerns, which is motivating these banks to just say, yeah, let's just wait on this. We're just going to outlaw this right now. We're probably going to experiment with it in the background, but we don't want employees using this at work. The other category of concerns though, and Chime I think is a good example of this, is companies saying, we don't feel comfortable doing that supervised fine-tuning of these models on top of models that are built by other companies. And the reason I think that they don't feel comfortable doing that, and the reason that I think Chime is wanting to start fresh by building their own model from the ground up and not taking advantage of any of the work that anyone else has done, is that they view the data that they would be adding to these models as proprietary.

And it's very early in the development of these things, but I think companies are already starting to think about, well, what data do we have that's proprietary and do we really want to share it? Right? Because anytime you use ChatGPT, just as a consumer playing around with it, you're actually helping to train the model. So you are giving OpenAI, which is the company behind ChatGPT, you're giving it data, you're helping it improve its model. Obviously me as an individual consumer, that doesn't really bother me. I don't think my intellectual property is so valuable that I shouldn't be giving it to this model. But if you're Chime or someone like that, you don't want all the things you are doing to these underlying large language models to then go back to those companies that control that because you're competing with yourself in a sense.

I think you're going to see more companies embarking on building their own versions of these models using the underlying technology or techniques, but keeping all of the data proprietary and not sharing it across companies.

Alan Kaplinsky:

What I'm wondering about, Alex, if you're going to build your own model you've got to hire your own experts in the area. And my guess is there is not a big reservoir of people that can be hired that know enough about the technology. Am I right?

Alex Johnson:

No, that's exactly right. There's two constraints to it. One is the people as you say. I can tell you that data scientists and engineers who are experts at generative AI are getting paid massive amounts of money by Google, by OpenAI, by Microsoft, by a number of other large companies that are investing in this. I'm sure the government is also investing in this technology. So those data scientists are extremely well paid, and it's not a big pool. And so your ability to recruit those folks is I think at least in the short term, while this knowledge is pretty specialized, is going to be pretty limited. The other thing is it takes a lot of computing power to build these models.

And so the other thing that you're going to find if you're embarking on this journey of trying to build your own one of these models is that even if you have the expertise, Microsoft, OpenAI, Google, they've invested huge amounts of money to collect

all of this data, to train these models, to then do the fine-tuning and supervised learning on top of that. That's a tremendous amount of computing power and expense. And even just operating something like ChatGPT and making it publicly available to a large base of users, it costs OpenAI a lot of money every time you ask ChatGPT a question. And that cost is going to be prohibitive I think to a lot of companies as well.

Alan Kaplinsky:

Let's take a look at the big picture and look down the road. And I'm wondering, let's look 10 years from now, what impacts will generative AI have on financial services 10 years from now? And overall, how significant is this technological breakthrough?

Alex Johnson:

To answer that second question first, it is a tremendously important breakthrough. I think that anyone who's studied the machine learning and AI space for a while, they view this generative AI revolution as a huge, huge leap forward in the field. And so I think that in terms of just how significant, broadly speaking it is, hugely significant. You referenced the idea that kids are not going to be writing their term papers themselves anymore. I can tell you my brother is an English teacher at a high school. Teachers are already talking about this is going to completely change the nature of instruction, the nature of take home assignments versus in-person learning. It's going to have that type of effect across a huge number of industries.

In financial services specifically, I think that as we've been talking about, I don't see these, at least within 10 years, these models being let loose to drive high stakes customer outcomes themselves with no human intervention. I think financial services just by the nature of how heavily regulated it is and just how much certainty consumers like to have about what's going on with their money, I don't think we're going to get to a point in 10 years where you see these generative AI models making underwriting decisions or recommending financial products without a human being involved. I don't see that happening. However, I do think that they will really have a huge impact on the nature of customer engagement inside financial services.

And so I think that 10 years from now, digital banking apps will all have a chat-based interface as a primary part of the experience of using them. And that interface will be engaging, it'll be personalized, it'll be responsive to you as an individual, and it'll feel a little bit like having a friend at a bank. In the same way that if you banked in the old days and you went into the branch every week or every couple of weeks, you got to know the teller and the teller got to know you. That's what the experience of using digital banking is going to be like 10 years from now, in large part because these generative AI models are going to make that interface between the customer and the bank much more personalized and engaging. I think that's one thing.

And I think the other thing is this technology is going to scrub a lot of the dull, monotonous back office work out of banking. I think that'll make banks more operationally efficient, and I think it'll make the work that gets done by human beings inside of banks generally higher stakes and more enjoyable.

Alan Kaplinsky:

Right, right. So moving away to the financial services for just a second and into the area of just an interested consumer, there maybe somebody on our broadcast who's wondering how could they use it today? Are there apps that are in the Apple Store that you can download?

Alex Johnson:

Yeah, there are a bunch of services that have become available. The place I would probably start would be OpenAI, which is the company that's at the forefront of all of these things. So they built ChatGPT, they built DALL-E, which is the image generation one. They have a number of other ones across different domains. I would definitely recommend signing up with OpenAI. You can create a free account and you can get access. And the access is somewhat limited based on availability. And as I said, these models are very expensive to run in production. There are plenty of times I log in to ChatGPT where it's not

available right now, please check back later. So you won't get perfectly responsive service, but you can get free access to these tools to just play with them.

And I think the thing I would recommend doing that's so important and is going to be like this will be the skill that gets taught to our kids right now, and then the future is like you have to have it on your resume. How good are you at prompting AI? Right? Because as I said, all of these things that AI can do with this generative AI revolution, it's all based on what you tell it, what the input is. And so prompting these AI tools, telling it what you want, and asking the question in the right way and refining your question based on what it gives you, that back and forth between you and the AI tool, that's a skill. And I can tell you having played around with a lot of these tools, it's a little different. It's not searching Google the way we've all trained ourselves to do. It's different.

You have to ask it in a different way. You have to then once it responds, you have to go, actually, I meant this, and you can have a conversation with it. So that prompting skill is something that I think everyone right now should be practicing because it's going to be a future skill that we all need.

Alan Kaplinsky:

I'm wondering, in the future, if you have Google Home or you rely on Siri or one of these other devices that can respond to you in a human voice, I take it that that is going to become a lot better, right?

Alex Johnson:

It will. Yeah.

Alan Kaplinsky:

Very often I'll get a response if I ask a question where it says, I can't help you or I don't know.

Alex Johnson:

Totally. Those tools which are a precursor to this customer experience that we're talking about, those were really built before generative AI came onto the scene. And so while they are an interesting thing to engage with as a consumer, as you said, they're not very smart behind the scenes. Their ability to understand requests is very limited to a certain set of requests. Those are going to get so much better. And so we're going to have those voice assistants on our phones and on our speakers at our house. They're going to get built into all the apps you use. There's a lot of really cool examples that have already been built. For example, there's one that's being used by computer scientists and engineers so that if you're coding, you can use this tool as a co-pilot to help you write code.

And a lot of the code that people write is constantly the same code used over and over. And so it's relatively easy for a machine to figure out what you need. It's going to revolutionize computer programming, it's going to revolutionize creative writing, just all of these different areas. At minimum you're going to have a really smart assistant sitting next to you that you can prompt for anything that you want to do. And in some cases that assistant is going to take over at least parts of your job entirely. And so that's the future we're headed towards.

Alan Kaplinsky:

Right. So how does somebody subscribe to your publication, Alex, Fintech Takes or listen to your podcast? What do they do?

Alex Johnson:

Absolutely. So if you just Google Fintech Takes, you'll find me and you can subscribe for free. I have a newsletter that comes out every Monday and Friday, and then podcasts on Wednesday and whatever podcast app you use, Apple, Spotify, whatever, just search for Fintech Takes and you'll find me.

Alan Kaplinsky:

Very good. Well, really want to thank you very much for joining us today, an absolutely fascinating discussion. A few months ago I did something on the Metaverse and how that might be used in the financial services world. I must say that generative AI, the prospects at least from my point of view, are much more exciting than the use of the Metaverse. My thanks to you, and perhaps as things develop and get refined, you might be a repeat guest on our show. But thank you again.

Alex Johnson:

Thank you so much.

Alan Kaplinsky:

Okay. So to make sure that you don't miss our future episodes, subscribe to our show on your favorite podcast platform, be it Apple Podcasts, Google, Spotify, or wherever you get your podcasts. Don't forget to check out our blog, consumerfinancemonitor.com for daily insights about the consumer finance industry. And if you have any questions or suggestions for our show, please email us at podcast, that's singular, podcast@ballardspahr.com. And stay tuned each Thursday for a new episode of our show. Thank you for listening, and have a good day.