Building Trust in Agentic AI: The Key to Success

CapTech Trends Podcast | Episode 47

Brian

Hi, and welcome to CapTech Trends. I'm Brian Bischoff. Today we'll talk about more of the business aspects of agentic and what it means for us in the industry. So looking forward to the conversation today. Today I have with me Liz McBride and Shane Sullivan, both are leaders in our AI space, specifically with all of our AI programs internally in CapTech. So welcome.

Shane

Thanks, Brian. Good to be back.

Liz

Thank you. It's good to be here.

Brian

So I guess as we get started on this, one of the foundational aspects of just any conversation in this topical area is can everybody kind of baseline on exactly what we're talking about? So agentic AI can mean a lot of different things probably to a lot of different people. I would love each of your perspectives on really what agentic AI is. How would you define it and what do you think are some of the key aspects of agentic?

Liz

To me, agentic AI is about the ability to focus on overall outcomes. And how we can use agents to make things more autonomous. So thinking about what is your business goal and then how can agentic AI help to automate some of those processes?

Shane

I'd agree with that generally. But for me agentic AI is really differentiated where past automations or different systems would be very strict in how you had to program each potential path. Whereas you can have an outcome that is known, you have a goal in mind and you can have the agents go through very complex, very in-depth things to get to that outcome. And that's really what separates agentic out for me from other technologies that we've used in the past with clients.

Liz

I think that complexity too Shane is the other thing that's so interesting to me since I focus on the people adoption aspect of it, is with complexity and with automation comes with it where it starts to change potentially people's roles and opportunities for growth and how people think. Because AI is able to assist and make decisions autonomously. And so then what is that human AI augmentation? What does that look like?

Shane

Absolutely. And one of the things we've been talking particularly in healthcare about is the fact that we want systems that can help bring time back to some very talented, very specialized people who need to do a lot of work. But it helps them make a decision. It doesn't make a decision for them. And that's key.

Brian

So each of you talked and I'll cherry pick a few words just because I think there's things that I heard that really help define what this is. One, you talked about automation. You also talked about being more outcome driven. And both of you used the word agent specifically. So I'll start maybe from the back side of that, but what is an agent from your perspective? We talk about agentic as a solution, but what's an agent in your mind?

Shane

In my mind, an agent gets very down to a detailed level. So it's doing a very specific task. It's a prompt, it has tools and it has knowledge. So it can access specialized knowledge to complete the task you give it. And then it can interact with other tools that allow it to be more than just your single shot prompting that you see with large language models.

Brian

Yeah, I think when I think about an agent, it's definitely narrowing into that smallest component of work that can be done. And I think you talk about automation. I think that the goal of that is really to be more efficient overall. But you also used more the outcome-driven approach. Is that a new philosophy? How do you rethink being more outcome-driven as opposed to maybe what we were

previously?

Shane

It's not necessarily a complete shift, but the realities that you can reach with agents are much more malleable. You can provide less structure and get the outcomes that you want. You might have an outcome desire in the past, but you had to get it programmed very specifically to get that outcome. And if it didn't fit within those guidelines, it had to go to a human. So agents can handle more complexity and more unknowns than past systems.

Brian

I liked your malleable word. I can't even barely get it out but I like that word. Because it definitely, when I think about designing systems from a technology perspective you typically have very specific scenarios that you're working towards. But that's more, I think... And also a shift in this. What I hear you say is that there's a shift in the way you think about the outcomes don't have to be explicitly discreet. It can be more malleable or more fungible or whatever the word.

Liz

You're not limited necessarily into where the data is, right? I mean, that's still an important aspect, but agentic AI can source information in various ways. And so it opens the door to kind of that product mindset of what are we trying to accomplish and not be limited by that finite task that we're trying to do. But really at the end of the day, what are we trying to get to? And then this is where the innovation gets exciting because if we start with what an organization's vision, goals, objectives are, right? You can not be so limited thinking of what this LLM component can do, but you're thinking of, "Wow, if we could redefine and be differentiators in our space." Agentic AI can open the doors to different ways of working. Yeah.

Brian

Let's move out of the abstract then, because you talked about specifically use cases matters and use cases. What are the use cases people should be thinking about where to start with this, right? I think using the word agentic, people think call center agents or travel agent type scenarios are the ones you always hear about. But what comes to your mind as far as the best use cases for solutions like this?

Shane

One of the areas that we've been spending a lot of time in at CapTech and with our clients is with healthcare. And a use case that's particularly resonant across the healthcare industry is the processing of claims. And so the specific example that we worked on and built out a POC for is taking claims for worker's comp and getting to the end of, "Is this an approved or not approved claim," based on very complex criteria. So that's a use case that would involve doctors and nurses and multitudes of different standards that are required to be evaluated as well as a patient's complex medical history whether they qualify for the treatment that's being requested. And so that's the level of complexity you can think about from a use case.

But we can get into how do you build something like that out? It's, "Is the claim valid?" Which has its own criteria. And we can have an agent that just evaluates the claim. And then we can have an agent that looks at a patient's medical history and the standards that are put out by the states to say whether or not they would qualify for the treatment that's being recommended. And this ultimately gets doctors and nurses who are doing the final approvals to the decisions faster. And they can link back to the source documents and the documentation to say, "Yes, for this reason, this is approvable or not." That's a lot more complex than, "I need to reset my password with a customer service agent."

Brian

Now that specific example though, that's not a new problem in healthcare. It's something that's been around for a while. And people have probably solved that problem a variety of ways. Why is this different now? Why is it all of a sudden that's a great use case? What's so unique about that or what's the opportunity in that that makes this a great use case?

Shane

For me what's unique about that is there's so much complexity in the standards and so much variance across what those standards are based on the different states that it's hard to program that. It's also hard to find talented people that know that broad information. So your ability to do it programmatically and to use agents to get doctors and nurses to a solution or to their decision more quickly is pretty invaluable.

Brian

So what's different now is there's these LLMs, these large bodies of knowledge that have this information.

Shane

All of these codes at once, you can update it. And when a code update comes out, the system knows automatically that these are the new standards that it needs to align to. There's no learning cycle. And that's really what's different is we can get all of the information in and updated and make sure that we're up to date as opposed to past systems which may have had a little bit more lag time.

Brian

That's a pretty big leap for people to get there, specifically in healthcare. That's a great example you described there. But also to take the leap and say, "I'm going to trust this system now that is very complex, has been hard to automate previously, and then we're going to leverage AI and agentic systems to solve this. And it's just going to work." How do we get over that hurdle?

Liz

Especially in the healthcare space, because it's high risk. So what we see in high risk industries like healthcare, transportation is it's even more important to establish that level of trust. And so people who have average AI literacy, meaning, "I sort of know what this is supposed to do. And I know its limitations in a way, but not really."

Brian

I think more people know limitations right now then know what's possible.

Liz

What's possible. So if they're in the middle of the road, what the research finds with high risk industry is they would be most likely to either mistrust the outcomes, or on the flip side they would overly trust. Does that make sense? Because it varies, especially low AI literacy would be like, "Oh, I don't know, I'll just..."

Brian

There's the skepticism or there's the blind trust. Both of those scenarios that you probably-

Liz

Exactly, neither are ideal when you are figuring out a code.

Shane

And the examples that I was talking about for healthcare, what we try and build in is the ability to allow the experts to link back to the decisions that the agents made. So for each action the agent takes, there's traceability back to how it got to that decision. And that's key. And so in the proof of concept that we built the agent would output a recommendation that a doctor or a nurse could then make a decision on, but it was key that they linked back to the documentation, the original medical file, the standards that put out by the state. Those are the ways that you build trust because if they were doing it by hand they would've just had to pull up that PDF, look through that medical file. They would've had to pull up the standards or know the standards in their head hopefully and be able to reference that. And so quick reference builds trust.

Liz

Explainable AI builds trust. And having multiple explanations as well. So going back to the healthcare scenario. So if you're using AI to look at a scan for results, what they found was if the AI provided multiple explanations then it instilled trust. Because to your point, it's showing them where it got the insights. But it's allowing for that doctor or nurse to also leverage their expertise to make the right call based on that. And what's even better, it's improving their own meaningful work. Because it makes them feel more confident in their own ability to decide on the right course of action or result. Does that make sense? It's like AI is actually helping them feel more confident because it's providing them with those various, "And here's how we came to that."

Brian

I think this is a fascinating use case. And we're diving into trust right now. I want to go one level deeper on trust. Because you've described the fact that the information explainability helps to provide trust. But when I think about trust in this use case. It's trust of the nurse and the doctor, but also patient trust. It's almost two different aspects depending on who the users of that information is, right? So do you treat those different populations differently when you're trying to build trust or you're building

adoption of these systems, or is it really just not that differential?

Liz

In both cases it's important that they understand where the sources are coming from. Just like if we were to just use ChatGPT today, I want to know where it got those sources. Same is true whether you're a consumer or an employee. So some of it is the same.

Brian

In my mind it's really what's different now than how we used to build systems. We used to build systems, we still had to consider trust whatever the technology was we were building, why is it different now? And when I start a new program next week that's going to be leveraging agentic technology, what do I need to consider now that's maybe different than the way we did things in the past from building that trust?

Shane

The building trust component with inside enterprise and with customers, when we're talking about those types of use cases. And in agentic AI specifically we often look at what are the source documents you're using today? What are the trusted sources that you have? And that's what we're going to enable your system to work from. It's not going to be accessing or leveraging other information. It's going to be leveraging the information that you know and you trust and you use on your day-to-day now.

And so I think that's where you start to establish trust is by creating these knowledge bases for the agents that are the knowledge bases and trusted sources of the organization already. And from a consumer perspective, I think it's allowing a consumer... If you're building something that's consumer facing, to be able to get the output based on the interaction, but then be able to go a level deeper and make sure that they can look at the original citation. So whether that's pulling something from the BBC or somewhere that's another reputable source, you're going to instill trust with the consumer by directly providing the source of the material. Whereas in the past, that was a lot more difficult to do. And you look back a couple years ago in some of the AI interfaces, getting to, "How did it come up with that?" Was very difficult because it was just a numbers machine of predicting the next word and creating a great output that seemed convincing, but sometimes was just hallucinating.

Liz

I think the sameness from an employee to consumer aspect is the usefulness of it. Is it a fit for the task at hand? So is it helping me? There's also a sense of well-being and benevolence. Is it intended for a positive purpose? Both employees and consumers alike care about that. So there's some of those aspects of it maybe not as tangible. But there's also because of agentic AI comes with it, the increased focus on the ethical aspects of it and bias. And so organizations need to make sure that they're considering clear guidelines for using these tools. So that heightens it. It also provides increased aspects of psychological safety because of the autonomy. So it's, "Will this replace my role?" And again, going back to that trust and is it making the right decisions without me having that degree of oversight. So there's various levels of aspects and flavors. Some are human AI, some are more autonomous than others. There's a lot of different opportunities, but it should fit the task at hand. I think that's really important.

Shane

I would say when we're talking about bias and decision-making, what's interesting about some of the agentic AI use cases is that you can perhaps look at what has been done historically by humans doing a certain process. And we'll stick with the example of different types of medical claims. There's always concern around bias of, "I'll approve Brian because of this reason. I'll not approve Liz because of that reason." And that's not good. But you can have great traceability of what decisions are made. And you can enforce standards to an agent in a way that is very consistent potentially compared to 15 different people processing claims different ways and them interpreting them with their human decision-making as opposed to agentic.

Brian

My hot take on this, I think as a society in general we've become very trusting of technology and what it can do in various places. But with generative AI and AI in general there's a skepticism that many people continue to have. And so coming into these solutions, whether it be you're an end user of an output of a particular agentic system, or you're secondary, you're the patient in this scenario and you're not actually interfacing with it, but you're getting the results of it somehow there's probably a healthy skepticism. And in order to get past that, you have to be thinking about trust and building trust and building these sorts of just explainability inside of these systems as a part of your programs in order for them to be successful. Because that trust that we naturally have in technology isn't immediately there with agentic solutions. At least not right now. That's my opinion on that.

So we spent a lot of time talking about trust, diving into that use case and talking about trust, talked about building trust with users, potential bias of systems. We need to work through. Potentially complicated technology that also, we haven't really talked about that, but there's a new aspect of technology that has to be figured out. Why would I even want to approach this? Why would I change the way that I want to solve problems? What benefits do we get?

Liz

Why wouldn't you?

Brian

What benefits do we get out of it, right? As a business, why am I going to invest in agentic versus something I've already done? What's the positives there?

Shane

Part of the reason I think agentic AI is so compelling to invest in is you can build a framework and an approach that is very modular. So you can start with the small set of agents that I talked about and build up to something that's very complex. And at the same time it's flexible in adapting to advances in AI, which are happening very quickly. So when you're thinking about something a year out, you know that you can take a new large language model that's a better performing one or a different one and put it into your system without a lot of rework. And so the level of technical debt is very different than other systems where you pour in a ton of money, a ton of effort, a ton of time to build something, and then it's very rigid and hard to revamp and update. You can make these updates that really change the behaviors and the performance quite materially with not as much development effort.

Brian

So we would spend a lot of time in previous programs defining those business rules, building out that logic and having engineers build that system from the ground up. What I'm hearing you say is that there's a different way to solve that problem from a technology perspective. And that provides a lot of flexibility for the future as if new opportunities come down the path from a technology, what's new with these large language models. Or your business changes and you want to adapt more quickly and it provides that sort of flexibility.

Shane

Exactly. You can stick with the same platform if you decide that your business process is going to change. To make that adjustment in the system is very flexible to do. It doesn't require retooling the entire thing and starting anew. It's making some adjustments with inside of the existing system. It's doing that rigorous testing from the evaluation set had from day one and making sure you're getting the same or better performance and then being able to deploy it. And those turnaround times on updates are very fast.

Liz

I think that makes the entry easier if you think about it. So I mean, with the rate of advancement of our tools and you don't have to wait until it's perfect. That's going to happen. There's going to be evolution. So to your point Shane that modular approach, it just allows organizations to go ahead and get started and not wait. But at the same time, it provides the ability to differentiate as well. Truly think about your business differently.

Brian

Because that's one way to think about it is you like, "Well, maybe I'll just wait for that next thing to come and just I'll jump on that bandwagon then."

Liz

Yeah, the lower threshold.

Brian

But really I think what we're seeing is let's start now, because the benefits really are there from both the flexibility of the systems that you're building, the capabilities that you're building from a business. And that's really only going to accelerate as we move forward.

Shane

It's more flexible. I think over time it's also getting more accessible and easier to use. If you look back five or so years, you would need very specialized skillsets. AI engineers and machine learning engineers

and those people are still terrifically talented and have great skillsets. But more and more of these tools are becoming accessible to at least start sketching out your ideas from everyday business users. And so to take business requirements and to build a POC takes a lot less development time and effort than before. And so you can really program your ROI and prove out the value of these systems with less development effort. And more of that conversation with the business of, "What is the true value you're trying to get out? What is the outcome you need?" And again, these are outcome-driven systems. So you take that outcome, you sketch it out with a few engineers and some business folks. And you can show that on a small scale it works and then make the decision to go further and modularize up to more agents. I don't know if I made it modularized. But yeah, we'll take it.

Liz

I like it.

Brian

Maybe quickly, what are some other great examples or example use cases that come to mind for you as potential applications of agentic that are maybe beyond the typical call center type thing that people immediately go to?

Shane

All right, yeah. Calling us out on going into the easy one. But I think some of the other use cases that get me really excited that we're seeing as high potential are in financial services. You can look at using an agent to do anomaly detection and evaluate against risk and compliance standards at very high scales and with very detailed levels. So before you might have had a limited number of models evaluating your overall risk exposure or your anomaly detection model. You could have 15 different agents looking at 15 different types of anomalies without having to put in kind of historically what the level of effort would be to build that type of 15 model version. And then in a similar vein, you think about high volumes of data coming in and having to process a lot of that. When you think about cybersecurity, the cybersecurity professional is receiving a ton of alerts and a ton of kind of things they need to consider. And so using agents to specifically parse through all of those type of inputs is another very high potential use case that I've seen a lot of opportunity with.

Brian

I love both of those examples. Because I think what I heard you say was that if we were to do that previously, you you're building these either large AI machine learning type models that have a variety of different inputs. And you have to be all encompassing in those sorts of models. This decomposes that to use the word you talked about earlier into smaller functions, smaller agents that allow you to do very discrete decisioning. And then that broader orchestration of those decisions is really where this agentic model comes in. I think that to me, that's what I took away from your conversation is what the possibility is.

Shane

Exactly. It's really bringing it down to you can build something that's very specialized and unique to the subset of patterns that you want to look at as opposed to a broad, generalized model that will catch some things. But this ability to modularize down to your exact pinpoint occasional anomaly, you can really customize and catch more things and reduce your overall enterprise risk. Which is always a good investment.

Liz

On the opposite side of risk, when we want to have fun and have engagement.

Brian

Talk about having fun.

Liz

In the sports space, in the hospitality space. How can these organizations think about what is the consumer fan experience? And so how can you think about all of the aspects of their overall journey and kind of bring it together and optimize it, leveraging agentic AI?

Brian

So we talked about a lot of different use cases. I think we've gotten your perspective on where we should be spending more time as these programs initiate building trust and making sure that we're thinking about these from more outcome-oriented scenarios. If I'm listening to this right now, I haven't done an agentic type of solution right now. Where do I start? How do I get this rolling inside of my

particular team, my business unit, my company? Whatever the case may be, where do I start?

Shane

Yeah. If I was a people leader I would think about what do I have my people doing that's taking up the most amount of their time where it's I have an outcome that I want them to reach. And it's time consuming, it's tedious, it's repetitive. And what I really care about is what is their analysis or their final decision. But to get to that final decision it takes a lot of time. You can then modularize down or kind of think about an agent for each of those steps. Whether it's going and doing some research, and then that research is relied upon to make a decision.

If you can just start with that one sub-component of your employees or your team's process, that's where I would start. I would start just thinking about, "What does my team spend their time doing? Okay, what do I want my team to be spending their time doing?" And then talk about what business value is and making sure that you're directing towards both ROI from a direct perspective and then a time savings perspective of making sure your employees are engaged in the most valuable things for your business.

Brian

To me, both of you hit on establishing really why are we doing this? What's the purpose? If you understand that purpose, whether it be, "My team spends a lot of time in one area," or, "My competitor is doing something that I can't do right now," or, "There's just an opportunity that I see in the marketplace. Maybe I can think differently." Understand the purpose for why you're tackling this. And then there's all kinds of things that flow out of that. How do we test this small implement, get small wins, those sorts of things? And then enabling your people also to think differently. I think those are a lot of different things that we've spoken about today.

Shane

And then making sure that it's not a solution in search of a problem. Instead, start with your problems. What are the problems or the challenges your business have? Because those are areas you're going to see ROI in. Don't take a solution and go try and find the problems. Just start with your problems. What are the challenges your business has? Then kind of work through that cycle of, "Okay, these are the challenges. Where is my data ready? Where are my people spending their time? Where do I want my people spending their time?" And that's a really good way to start mapping out your journey.

Brian

So we spent a lot of time talking about what's involved with from a business perspective thinking about these agentic systems and how to decompose them and solve those. But there's an employee aspect of this too. What do we expect? What should we be expecting of our peers, our employers and employees really in this scenario about agentic systems? What do they need to do in order to be successful?

Liz

It's important going in that we think of it both from a technology and a human and thinking of that evenly. What I mean is I think in transformations in the past, it was technology led. It can no longer be technology led. We have to think about the whole organizational aspect in the human and the technology together. So that being said, thinking about what does the future state going to look like? Again, we don't have to have all the answers. It could be incrementally, but using tools such as job crafting in order to ensure that the tasks of the human is optimized at the same time that AI is optimized as well. Does that make sense? So how can organizations really leverage both in order to get the value that they're after?

And so what job crafting does is it breaks down the role of how the employee and AI will be interacting from a task perspective, from a relational perspective. So how am I going to now start collaborating? I may collaborate with AI. AI may be a supervisor. I may supervise now AI. I mean, there's so many different flavors of what's possible. And so how can we help that collaboration? And then the third is cognitive tasking, where it's about, "What is possible now? What is my purpose? What is meaningful to me?" Working with this tool, because with agentic AI, it may change the purpose of a role in general and the opportunity to think more strategically, to be more creative potentially.

Brian

I think it may change the purpose of the role, but the reality is it probably will change the person interacting with that responsibility.

Liz

And that's great. It is a positive thing. And it's really important to make sure that employee is part of that process. So when I say job crafting, the key there is it's employee driven. It still leverages leaders

in HR certainly to create the parameters. But it's the employee testing the AI outputs the tools. Again, understanding the explainability transparency. And then they're using that to future design how they can work with AI. And by doing that you are setting them up for success. And you're also getting the output, and probably even more so. And so employees who leverage job crafting when it comes to agentic AI are more inclined to adopt it, more inclined to trust the output. It's early days, but it's showing strong indicators of job satisfaction and increased job performance. And so that has to be coupled with making sure that we have comprehensive AI literacy programs that's not just about the AI algorithms.

Brian

Yeah, it's not just about the output of what the AI creates. There's a huge aspect of... I love this job crafting concept. I think what I'm hearing is you just assume that these roles are going to change somehow. And we need to make sure we're preparing the individuals in those roles to adapt their model to support this new agentic system.

Liz

It's like, "Hi, Brian. You get to use this tool. So guess what? You get to be creative." "Well, maybe I don't naturally want to be creative. And maybe I enjoyed what I was doing yesterday." So getting that aspect of it. And then making sure that we're providing ongoing continuous learning from a technology perspective, from a responsible AI perspective. So how do we ensure that we are thinking about ethical and mitigating bias and all of those things. And then thirdly, what is that human-centric aspect of opportunity? And that includes soft skills. So especially with agentic AI with the ability to reshape jobs, it will require humans to increase their emotional intelligence levels by six times what we had to previously.

Brian

That's a tall as potentially, right?

Liz

Six times, yes. And also we say it unlocks innovation and creativity, which is exciting. The reason that's important for employees--and all of us really sitting at this table--is AI cannot display empathy like humans can. It can have context eventually and so on and so forth. But especially things like

compassion, AI can't be compassionate. It can't. I'm just not allowing that.

Brian

Fair enough.

Liz

We say AI can be innovative, it can. So can humans. But we don't even know the human brain, how we are creative. We haven't really even mapped our own brain, so we can be differentiators in that space. So I think this is a great opportunity. I'm excited about what's possible from the human side. I just want to make sure that the human is in the loop in the outcomes. But I really want the human in the loop from the beginning of just making sure that they're part of the process, part of the solution. That we're designing it in a way that is adding value. But honestly, is just meaningful.

Brian

I think that's a good place to close in. Liz, I appreciate your time and thoughts on this topic as well as Shane. I think this is just a continuing space that will evolve, but it's great to get perspectives on how to really get started with this, because it's certainly an opportunity that's not going to change in the near future for us. It's the direction we're going, so we got to be ready to jump on that change. Appreciate the time. Thank you.

Shane

Thanks, Brian.

disclaims any and all liability or responsibility for any direct, indirect, incidental, or any other damages arising out of any use of or reference to reliance on or inability to use this podcast or the information presented in it.