



Agentic AI: Solving Complex Problems with Specialized, Autonomous Agents

CapTech Trends Podcast | Episode 48

Brian

Hello, and welcome back to CapTech Trends. I'm your host today, Brian Bischoff. I'm excited to continue our conversations in the AI space that's continuing to evolve, and specifically today around agentic AI. It's been getting a lot of attention in the marketplace. Today I have Jason Dodds and Kevin Vaughan with us for this podcast.

Kevin, you've been a frequent participant in the podcast. Jason, this is your first time, but welcome.

Jason

Thank you.

Brian

Looking forward to the conversation today.

Jason

Looking forward to being here.

Brian

Just a little bit of background for context, Jason, your background's more aligned to a long history with data, data analysis sorts of engagements and just your background in general. Kevin, very broad mix of both innovation and technology, and certainly innovation in the AI space. Really looking forward to both of those perspectives shining through today.

All right, well, let's get started because agentic AI is really something that's been, I feel like, getting more and more traction over the last year, I think really the last six months. It's hard to flip on some sort of technology news or some sort of product launch and see that they're not talking about agentic. Maybe if you could start really with just each of your perspectives on really what agentic AI is and how it's really different from things we've seen in the past in this space.

Kevin

Sure. I'll take it first. I think that the evolution that we've seen of AI over the past, call it five years or

so, has really shown what we can do with generative models. If we look at generative AI, large language models, ChatGPT as it first was coming out, what you see is a lot of knowledge model. Things that understand the world, things that you're able to speak to and get some of that knowledge back out, but it's a very transactional, "I'm going to do this and this is the answer I get back."

What we see with agentic AI is to start to use some of this generative AI power for reasoning and be able to then construct multistep plans to accomplish a particular objective. And then use tools to actually act and have an impact out in the world. Whether that tool is to go and retrieve information or to actually take action. Think about customer service bots that you may have been able to answer questions that you would have about your policies or whatever before, and now you can already go ahead and start auto-renewal processes and do that same thing conversationally.

While at the same time, from a technical perspective, really the conversation structures that you had to set up before were very complex. Lots of I need to define all of the landscape of what's going to happen. Now we're more back to I need to define the objectives that I want and my agentic AI system can use the tools and the data that it's been empowered with to be able to accomplish those for me within some degree of guardrails.

Jason

I think for me, the concept of agentic AI really comes down specialization. Kevin, you just talked a lot about generalists becoming specialists and how do you build a network. But how do you create someone that is extremely specialized and performing one thing over, and over, and over again? If you look at that in today's workforce, you spend a lot of money to go find somebody that's really talented in one thing in particular. We can take a lot of that capability that somebody would have as a specialist, in my space, somebody that knows how to run an event. You're going to hire someone at an event that knows how to run a food and beverage center. How do you go train an AI agent to understand or series of agents to understand how to run a food and beverage center that is part of a bigger whole of running an event? So that they can start working together, so that we can truly start specializing each of these agents to focus on one thing, know it better than anything else in your network to be able to surface insights really fast.

I think for us in the sports and entertainment space, it's how do you surface insights faster so that you can give information at a higher level to the people making the live event take place. How do you surface that information faster so that humans can still make decisions?

Brian

Jason, you introduced at least one new concept in your opening description on what agentic is and you talked about an agent. I don't want to make any assumptions that people understand that because I know that, certainly, who you talk to can have a different perspective on what really an agent. What is an agent, and how does that contribute to this whole agentic conversation?

Jason

When I talk to clients, I really think of an agent as someone who is specialized in doing one thing really, really well. Whether that is go and get data from the open web, go get data from a system that's internal, or surface an insight. We're going to go create an agent that does one thing really, really well. They're, in my mind, very task-oriented, where they're going to go and execute against one thing. If they don't need to execute, they're not going to be part of the system.

They are, same as you'd called into a call center agent, that call center agent, if you're calling about your credit card, is going to help you resolve a credit card dispute. Same idea for AI is you call in and you're accessing this piece of a network that does one thing really well, and can hand off the relevant pieces to the rest of their network to figure out how to solve problems for you.

Kevin

I think that problem solving as well, tying onto that, is the dynamic nature of it. And being able to figure out what would be applicable, what data do I need to actually go out and get to do this instead of having to pull everything. And do so autonomously. I think that's the big word everybody throws out is autonomy and being able to do that on the fly adjustments to whatever the scenario is.

Brian

I want to get really specific here because I feel like we're at a point where it's very easy to toss around some of these terms. People could say "we're building an agentic system," or "I built an agent that does this," when in reality maybe it's just calling an API of a system and retrieving from it. How do we not abuse that term? Does it have to have some sort of model built into it? Does it have to use generative? What's your perspective on really what qualifies as an agent?

Kevin

Agentic I think is a nice broader term that allows some degree of flexibility from how does this thing actually run, how am I actually oriented. But I think naming and versioning are two of the biggest challenges that we have right now in the AI space.

Jason

I don't know that an agent is much different than API, in that an API really is just getting data. I think the agentic portion of it being an agent means you can do more than just get and retrieve it. You might do one step and format it, and pass it off to the next piece of your system. But you're doing more than just doing retrieval, push and pull from the system. I think to be an agent, you have to do a little bit more than that.

Brian

Yeah, when I look at it, there's a degree of variability in these agents with some of the language models or other capabilities in AI that I think provide some flexibility. It's not just "I'm going to get these three inputs, I'm going to do this one output." There's some concept of variability in there too, right?

Jason

I think when you start building these networks of agents, I know we recently built one in sports that were looking at agents that were using bespoke AI models, even different generative models. Each agent along this entire network, some may use model A, some may use model B. How do they work together to get to the best outcome? Some models are better at, for instance, doing math, and some are better at creating generative outcomes. How do we think about what each one is doing along the way to create the best outcome you're looking to drive in that specific scenario? I think that's where agents become more powerful than just hitting an endpoint on an API, is you can start doing more with the data than just straight retrieval.

Kevin

I think the orchestration there that you're calling out is one of the big pieces. I can hire somebody that can do some sous chef prep work. I can hire somebody that's going to do fraud analysis. In both cases, I'm just hiring somebody. If we talk about agents as employees, that description is going to be equally applicable. The level of effort that goes into each of those tasks, the level of training, knowledge is dramatically different. Even just the type and the domain.

I think it's fair to have some confusion on the terminology because it is general terminology. When you think about an agent being, as you said, very task-specific, very honed in, it does feel like an API. I think from a technical perspective, I've had conversations where laying out, "Okay, here's an architecture that you can use for an agent, here's how you can deploy this." They say, "Oh, this is service-oriented architecture. This is circa 2000." The reality is, yeah, if you're looking at it at an individual component, it very much is. Why not fall back on decades of engineering experience that shows how to put good systems into production at scale?

But when you put them together and you start having their objectives overlap, and you start, "This person's an expert at using these tools. This person's an expert at analyzing and insights. This person's at expert at talking about particular things." Now you can solve a problem that you never explicitly actually said I need to solve. You're using emergent behavior coming out of the combination of these agents to accomplish something that I think is really where people expect magic out of AI. I think that's where it starts to become magic.

Most people aren't looking at that kind of thing yet. They're maybe disappointed at this moment, but the work you're doing in sports, we've done similar multi-agent systems in logistics that we've talked about before. It's the confluence of these different agents coming together that really I think makes that difference.

Jason

I think where you, or Brian you started the conversation today, is my background's a lot in data. Kevin, yours is a lot on building systems. I think what AI enables us to do is I don't have a background per se in going and developing APIs or getting data from endpoints the same way that you may when you're developing software. But what an agent allows me to do is go do all of that with English language. I can go and accelerate my ability to access that information and do something with it that I otherwise wouldn't have to build that network. I can go from I have an idea to I have a full proof of concept at the end of the day. Within that eight-hour workday, I can go from end-to-end by myself without having to work with a team of six people to develop that prototype.

Now, not saying that the team of six people wouldn't make a better product at the end of the day, but it at least helps you start thinking through what is possible faster than we did before.

Brian

Yeah. I think I put each of you in your own little corners, one data, one systems at the very beginning. Not intentionally to divide that, but I think we are seeing a very strong blend and overlap there in that Venn diagram of skillsets in this space.

Kevin, you mentioned a couple of things in your description so far. You've talked about reasoning, the capability of these models to do reasoning. We've talked about orchestration. When we talk about agentic, it's the concept of orchestrating a series of different tasks or agents to do things. When we talk about it from a broader CapTech Trends, in a previous podcast we've talked about that being part of an automation story, and how companies are trying to build as much automation as possible.

Where are we right now? What's possible now? I love talking to you, Kevin, because you talk about really what's the next phase and what we should be working towards. What's possible right now? What's the current landscape with these things? What sort of use cases or maybe examples that you've seen that are working right now?

Kevin

What is possible right now is more than enough for more people's needs. We talk about AGI and ASI, and things of what is the future going to go, and what does universal basic income look like in these more futuristic looking perspectives.

For tasks that I've got, especially things that I've got a lot of training material on that I would be able to use to help in the development of agents for any kind of roles where, think about high turnover kinds of scenarios. You've got a lot of documentation because you're bringing somebody on and you need them running very quickly. Again, going back to that employee hiring perspective, how long does it take to get the expertise needed to contribute at that particular level? There are some state-of-the-art bleeding things where research papers are being written, but by-and-large, you may see PhD-level knowledge being expressed through these world models, but you're not seeing PhD-level thought yet.

I tend to think of it as what's realistic. What can we do? I have no questions whatsoever. You're talking about a use case, and I look at that and I say, "Yeah, we can build that, no problem." I can at least correlate it to several things that we've done. I just know that that's going to be right down the center. I think of things that are maybe more ambiguous where there are some areas that we're not exactly sure how the systems might be able to handle particular levels of reasoning, or particular types of complexity of scenarios. Think about fraud detection, where I can try and pull everything in to a traditional ML model and be able to see is this transaction fraudulent or not. But I can also stick an

agent on that and say, "Yes, I can do that for this one transaction." But I can also look for trends and patterns that are happening of suspicious activities, maybe across multiple channels, things that you would never have been able to train a traditional model on, and be able to retrieve all of that data, and do that kind of inference.

Brian

There's a lot of expectations, or a lot of things that people probably need to do to start on this journey. You talked about just baseline knowledge, understanding what's possible so that we can get to a point of even just envisioning what to do with the data you have or what to do with the opportunity that's in front of you. You need to have some creativity. Or understand what's possible, and then have some creativity to march down that path. Governance of data, governance of the system, governance of the processes as you implement them also need to be there.

What else? Someone as an engineer on a team, as a data analyst on a team, as a team leader of these programs that are starting to explore going down this path, what do they need to be concerned with?

Jason

You have to start somewhere and you have to be willing to accept change. I think there are two guarantees in life, and I'll add a third. Everyone says death and taxes are the two guarantees. I'll add a third that says technology will change. If you're working with technology, it's going to change. If you're not going to start now, you're never going to get started. You have to start somewhere, and at least dip your toe into it, and use an open model on the web. If you want to go a little bit further and you say, "Hey, I like the results," that's when you start paying for access to a private model. You just take it little bit by little bit, and know that the same as you'd have to update your instance of Salesforce when they upgrade from version one to version two, you're going to have to do the same with your AI models. It's just going to be easier. There still is work to be done, but it will be easier because they're trying to make it as simple as possible.

Brian

In this new agentic model, what are some of the best practices from a design of systems that you need to consider that maybe is different than where it has been in the last five years?

Jason

Best practice would be start small and add. Don't try to solve everything all at once. Solve one problem, and then start layering the problems on.

Call me a simpleton, but I like to think of things as Legos. You can't build a wall of Legos without starting from the first brick. How do I identify what the first brick is? Let's go get that brick and then start layering on brick, after brick, after brick until I build my wall, or build my car, or build my sailboat. I think there's a Lego model now for the Eiffel Tower that's five-feet tall. You got to start somewhere. You can't start at the top, you got to start at the bottom. What are the base components I need for this to work? The same as you'd go build an architectural system 30, 40 years ago. How do you design a system from the components that you need to know work 100% of the time, and then start adding capability on top of that one layer at a time.

Brian

Well, I think that's an important aspect because what we're seeing, certainly when I'm in my conversations with clients, seeing is that people want to see value out of this as quickly as possible, too. They don't want to invest a year-and-a-half worth of time to get to something. They want to see how quickly we can get to proving this is really going to work and provide the value that we're anticipating. Those smaller building blocks I think is a great venue to do that as well.

Kevin

These small blocks coming together. Yeah, I've got to start with one, but you don't ever look at one and go, "Yeah, that's going to be the Eiffel Tower. This is going to be great." It takes a while for that vision to fully come into place.

Jason

I think you have to start with a vision of what you want to do, but you've got to start at the foundation. You can't start at the top and work down if you're building the Eiffel Tower, unless they've magically invented antigravity high bars.

Kevin

Next year.

Jason

Next year? Okay, great. Next year. But I think you've got to start somewhere. I think that's the hard thing to realize. Everyone wants instant value out of AI. I've heard that every day for the past six months. But the reality is if you're working with data and large systems, getting value day one is really, really, really hard. You've got to start small and build on it.

I always try to say, "Let's go find one dataset that you're interested in and let's see what AI can do with it that you may not have thought is possible. Let's show you that ROI with one dataset on day one." Then say, "Imagine if we did this with the next 500 datasets you had, and then built a system that could work across all 500 of those. What's the power that you guys could unlock with your business?" In my case, a lot of sports leagues. What's the power you could unlock with your fan by bringing together these 500 seemingly disparate data sources? If they were all working with a very, very tuned agent to understand every aspect of that dataset to create value back at the end of the day for the people that are consuming your brand.

Brian

You have to have some creativity about how to pull all this together. Then to your point, you got to start. That's why you can't wait for these things to evolve and mature a little bit. I think they're at a mature state now to do these valuable use cases.

A couple final questions here, just as we wrap up. One, coming back to the way I put each of you in your corners. Jason, from a data background, data perspective, Kevin, from a systems, integration engineering side of things. If you were speaking to somebody in that corner, how would you tell them to embrace this and get started? Where would you start to I guess really get the knowledge and start to test out what can happen inside of agentic systems?

Jason

I'd start by asking a simple question of, "Here's what I'm looking to do. Can you help me come up with an idea for how to solve it?" I think what's really great about these generative models is if you give it an idea, it can help you formulate an idea.

Recently, in the past two weeks, I had started with the idea for something else that I think we can do with one of our clients in sports. I was like, "Well, this probably needs a business plan. A business plan

usually takes about a month to write. I wonder if I could do that faster if I used AI?" The answer is yes. You just have to ask it questions and have a conversation with it. What's really cool about AI is it learns you, it learns how you react to things. These engineers that have built these large systems are really smart and they've done a really good job. Credit to them. But the outcome is you're going to be able to have more of a conversation with somebody that it's not just talking to yourself in a vacuum. You can ask AI to be critical of your thoughts. You can get started with a new idea. How do I refine it? How do I come up with a wire frame for something that I don't have?

From the data side, I use AI to help me generate the visualizations of an interactive system that I'd work with someone like Kevin to build. To say, "Here's what I think this could look like. Kevin, can you validate that this is worth building or this is the right way to think about this?" Let me give you something to look at rather than just says, "Here's this idea," and have some words on a paper. I can give you something to actually tinker with that is low function, but it at least gets you started.

Kevin

I think what you're describing there is use it in your actual tasks. Don't even try to think about how do I build a system with it. If you're not using it to begin with, then how do you expect to be able to build a system that can use it as well? Getting that practice in. Along the lines of helping you think through, I think removing white spaces is basically what I think of it. I never start with a blank canvas anymore. If I'm getting ready to do a thing and I don't have any thoughts about it yet, I'm leaning to AI to begin with. And say, "Here, let me see what a skeleton possibly looks like around this. What are the bones that I would need to be looking at or thinking about," as I go to talk, or design, or whatever it is that I'm doing.

But being able to then also contextualize that to me. Yes, it can implement memory and it can be able to recall things about me, and all of this. But I can also just say, "Here are the things that I know. This is what I'm trying to do. This is what I'm trying to work through. Relate this back, help me correlate." Take this T skillset that you get, and then just immediately be able to push a button to raise up any one of those lower data points up to a skillset that you need.

Brian

One of the things I took away from both of your answers to that was you have your baseline, you have your knowledge base of how you've built your entire career and everything you understand. But leverage AI to support how you take this to the next level. I think that's a mindset for anything in the

era we're in right now, but certainly as you're evolving and trying to learn about agentic systems.

Jason

Expand your capabilities and be able to have conversations with people you didn't think you could have a conversation with before. I could now have a conversation about the systems that Kevin's using at a level that I definitely couldn't a year ago because I've been able to actually see it come to life through prototyping, and just asking questions about how does this work, why does this work, and what's the general thought behind how this works. We need to think about how we can use the new tools that are evolving as the world evolves around us to help our clients be successful. I think AI is on the forefront of how we can do that for our clients.

The beauty of AI there is I don't have to pick up the phone and call Kevin. Not saying that I wouldn't like to do that, but he has other things to do. He can't always answer those questions. It gives you that very quick feedback on how you can think about getting started. I've heard this many times over in life is doing the same thing over and over again and expecting a different result is the definition of insanity.

Brian

Yes.

Jason

If all of us take our jobs today and just continue to do them tomorrow, and the next day, and the next day exactly as we are today, we're doing no good to ourselves, we're doing no good for our clients. We need to think about how we can use the new tools that are evolving as the world evolves around us to help our clients be successful. I think AI is on the forefront of how we can do that for our clients.

Brian

In 30 seconds or less, you've been doing this for a while. You've been exploring this space, exploring agentic as things have evolved over the last year, 18 months, maybe even longer. What are you learning about right now? What's next on your opportunity space?

Kevin

I tend to try to look at the periphery. You said earlier looking towards what's coming next and what the unrealistic world looks like. I find it helpful to know what they're not capable of yet and how you can mitigate against those things. If I'm being brought use cases, I want to understand what's possible. I want to understand what are the challenges going to be, and I want to think through how do we shape the solution then to either avoid those, or just acknowledge this is not really something that is realistic. In the same way it's operating in your circle of competence. I know what I know, and I know as soon as I start to get a little bit out over my skis, I know to raise my hand and say, "Let me ask some people and get some advice here."

Brian

Jason, what are you looking at right now?

Jason

Definitely more of the personalization and contextualization I think is where I'm focused. We've been, in the sports entertainment space, really focusing on the concept of the best experiences that we have are produced experiences. How can we give the people that are producing experiences for us, whether it's at a sporting event, whether it's an entertainment venue, how can we give them information to make my experience, your experience, Kevin's experience more personal and contextual to us faster? We have access to more of that, AI can distill it down into what is relevant to each of us. How can we put that in front of the people that are going to make it happen faster so that all of us walk away feeling we've gotten the most value and personalized experience for our dollar that we spent?

Brian

Yeah, I think I go back to a comment each of you made earlier, previously, about a lot of energy to build these large models. If you just decompose that into various pieces and figure out all the different aspects of what a personalized experience might look like for Jason, to be able to solve that individual components, and then have an agentic solution that can compose all that and make variable decisions depending on the outputs of those agents. I think that's a very compelling use case.

I want to ask a question here around each of you are very well versed on what's possible here. You understand the technology, you understand certain applications of it. But there's people you work with on a regular basis in leadership positions or decision making positions that maybe don't understand what's possible. How do you go about describing what's possible, convincing them of what's next, and

where we are in this whole ecosystem?

Jason

I've definitely started from a place of it's easier to show than tell. I end up building a lot of demos for clients, for people in the C-suite positions that are looking for, "Why should I invest in this? It sounds really expensive and doesn't sound like it's going to give me any significant longterm value more than what I've already got people doing." Let me go build you a demonstration of what it can do, and I can do that in three days, and you have a team of 200 people doing this full-time. If I can get to any semblance of what these people are doing with data I can find on the open web in three days, we should at least have another conversation.

Brian

That's an extremely powerful ... If you can turn around these proof of concepts in that short period of time, then that's an extremely powerful position to be in to demonstrate these sorts of things. I think that's a key takeaway for me certainly. When you're thinking about these things, it doesn't take a long period of time to build these, certainly to demonstrate them. You've got some governance and you have some things you got to figure out, but to prove it's possible, it doesn't take a whole lot.

Kevin

Yeah, I think there's a rub there that we are always constantly dancing the line between art of the possible and the reality of implementation. Very frequently, we'll build out POCs because seeing is believing, definitely better to show than tell, and we see that time and time again with all of our conversations. If you could show somebody, art of the possible definitely opens their eyes to, "Well, I can even at least talk about what might be realistic, even if that's not right." A lot of the times, those POCs are a golden path. I'm picking the right data, or I'm cherry-picking the problem that I'm going to be exposing this to. Not really necessarily painting a full picture of what this would take.

I go to go into production, I actually have to have an entirely different infrastructure in place, versus running it locally or running it in some siloed environment. I've got to have all different security things in place, and work with a lot of teams to get all of that. All of that data actually had real governance, or ideally it had real governance behind it, so where is that really supposed to be coming from? Is it fresh? How do I know it's good?

Brian

There's a lot of things that are part of the [inaudible 00:26:47]

Kevin

A whole lot of things.

Brian

... process you got to figure out.

Kevin

Right.

Brian

I think we were talking earlier, just how do you capture the imagination of what's possible? Often times, I hear pushback or feedback around what I guess I'll call the inertia to get started. Overcoming that inertia to get started because people are fearful. "My data house is in a mess, it's not in a great place. It's just going to take too much." How do you respond to that?

Kevin

They already started. If they're working for you and you're not giving them AI capabilities, they've found it on their own. That's just what the usage data shows, that's what all the surveys show. It's a giant shadow IT world that's happening out there and it's happening without the concerns that we've talked through around governance, and around verifiability, and around traceability, and responsible AI. You've got to start because you're already playing catch up. You really want to make sure that it's being used with your business in the ways that you want it to be used.

Another thing that you can do is to use agentic AI to help you. We've done some work using AI to help ingest your data in different ways, and help you form that data fabric, and really get those insights, help you curate that metadata so that you are able to interact with it better.

Brian

Well, thank you guys very much. I appreciate each of your perspectives. I know that we'll probably have some more of these conversations as we go forward in the future. Thanks for the time today.

Kevin

Yeah. Thank you.

Jason

Thank you.

The entire contents and design of this podcast are the property of CapTech or used by CapTech with permission, and are protected under US and international copyright and trademark laws. Users of this podcast may save and use information contained in it only for personal or other non-commercial educational purposes. No other use of this podcast may be made without CapTech's prior written permission. CapTech makes no warranty, guarantee, or representation as to the accuracy or sufficiency of the information featured in this podcast. The information, opinions, and recommendations presented in it are for general information only and any reliance on the information provided in it is done at your own risk. CapTech makes no warranty that this podcast or the server that makes it available is free of viruses, worms, or other elements or codes that manifest contaminating or destructive properties. CapTech expressly 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.