Houston Trend Talks

Navigating the AI Revolution: Insights on Implementation, Cybersecurity, and Operations

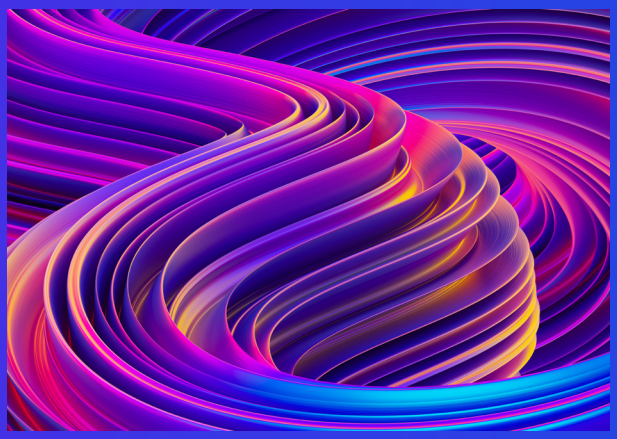

Josh Galvan (00:0):

Hi, everybody. Thank you for tuning in for another episode of the Houston Trend talks podcast. I'm, Josh Galvan, a KPMG Partner based in Houston, and I'm excited to be hosting a discussion today on the amazing world of Artificial Intelligence or AI. I have two KPMG leaders joining me for the discussion today. Anu Puvvada who leads KPMG studio and Jason Howard-Grau a cyber security principal in the firm. Together we're going to go on a tour of where AI is and how it is evolving in industry, how companies are using it or how they're approaching using it for potential benefits to the business. AI is arguably the most talked about technology development in recent times, referring to the sophisticated way in which AI solutions use data algorithms to create new and original content. AI as we all have seen and heard, is gaining much traction both in the business world and in our personal lives. And in the business world, particularly for its potential to improve and transform businesses, decision making, business execution, and of course, business results. With Anu and Jason's help today, we get to explore the areas of AI solutioning cyber security and operations. So, I hope you're ready. It's going to be interesting. We're going to go ahead and make a start by hearing from Anu on the areas of experimentation and implementation of AI. Anu, tell us about some of the areas where AI is being implemented and the partnerships that companies are forging to leverage its potential if you could, to get us kicked off here.

Anu Puvvada (1:23):

Thank you, Josh, it's nice being here today. Well, there's definitely been an AI revolution taking place. PitchBook had nearly 700 AI startups created in the last 5 years. If that's not a revolution, I don't know what is. The buzz word right now is Generative AI but it's important to remember there's many different forms of AI out there. There's machine learning, natural language processing and Generative AI and companies will need to pull together all different types of technology to activate their entire workflow using artificial intelligence. When I think about AI, I really think about it as a stack, there's applications that are built on the AI layer, ones that integrate into a company's main business model. And these are really applications that have emerged to be point solutions that fit certain problems or needs in a certain area. The big one right now is marketing and content creation, because that's one of the main use cases for generative AI we're seeing quite a bit of market movement in this area. There're companies like

Character AI, Jasper really solving for the marketing function. Another area is legal. We're seeing companies like Case Text and Patent PAL ones that really help the legal function operate better and faster. So those are really applications that are point specific function, specific solutions. There're also areas like Midjourney, which is another company that makes imagery that fits any industry or function. And then the next way I think about it is, in the stack are large language models which many of these apps are built on. So that's what we know as Open AI's Chat GPT, Anthropic's Claude, Google's Bard, Stability and depending on how these large language models are trained, they can give different answers. For example, Claude is built on constitutional AI inflections built with emotional intelligence. So how those pair with the use case are important. And then you have infrastructure companies. Those are the software and the hardware companies that really are the compute of AI, the Amazons and Microsoft's, the Googles and Nvidia. Computes are going be a challenge coming forward and we're already seeing it emerge. There's not enough compute in the world right now for all of the use that's emerging. And the last area, I would say, is we've been living kind of the last decade where companies have been collecting all of this data, but very few have really viable business models or revenue from capitalizing on that data. And what's powerful about generative AI is really the ability to be able to leverage an already trained model and be able to accelerate certain business outcomes. So, the question that companies should really be asking themselves is, how does this allow my organization to optimize my business? And how do I use this to create new businesses for my company?

Josh Galvan (4:37):

Wow! Thanks. Anu it's really awesome to be hearing this from you and imagining the potential transformative impact AI could bring to companies. In fact, the impact it is bringing to companies. And what's interesting about it to me is there's an entirely new vernacular for all of us to learn in not just technology, or technology enabled business processes, but maybe business entirely. And what's interesting about that vernacular to me, at least, as I'm hearing you is not surprisingly a lot of that vernacular is human like, you talked about emotional, you talked about constitutionally, about some other things as it relates AI solutions or AI purposes, and I guess, given it is artificial intelligence, that vernacular would reasonably verge on the human. So, it's really interesting to me and I'm aware, through a number of my discussions with my clients, that a lot of these well-known big tech outfits like you mentioned are right in the middle of all these collaborations. You talked about compute power and a shortage, and frankly, what I understand to be a long line of people in waiting to get their access to machines to do some of the engineering, the experimentation, and working on their own solutions. In your view, responding to that, are there any standout examples or really good success stories on how partnering actually can be best engaged or executed to actually produce a solid AI outcome. And what does that partnership model, if you will.

Anu Puvvada (6:15):

I think it really depends on the situation. Right now, we're seeing big players emerge, we're seeing startups emerge, and we're seeing companies make their own solutions using their own differentiated data. And I think whatever partnership a company goes into depends on which way they go. Companies need to really think about taking a two-pronged approach here. The first is around digital transformation and reimagining ways of working for themselves through AI adoption. And we're seeing companies play with the Microsoft's, the Googles, the Amazons of the world. But we're also seeing them work closely with startups in this space with those point solution. And the second prong of this is how does a company use its own data corpus and its own differentiated data to create new forms of revenue for themselves. And I think it's important

to focus on both areas equally and then the partner arrangements that come out of those is going to depend on, whether is it the first, or is it the second? We're definitely seeing companies race towards implementing Gen AI and taking an IT based approach. And I think it's really important to get this into the hands of people to see the use cases that emerge out of it. But the people and change management aspect of it is critical. At KPMG pretty much everyone has access to be able to use Generative AI. We're working with many startups in the space as well and implementing it into our operations, but also going to market with these companies. When we deploy generative AI at our company, we put certain guard rails on it and want to ensure that people would optimize the use in their day-to-day job.

Anu Puvvada (8:02):

It's also important to give people the right training and critical thinking skills to ask the right prompts and questions from this technology that's emerged. And the question to ask is, how do we get humans and AI to work together? And what does that model look like to amplify value, not just for the employer, but also for the employee and the partnerships that companies make, whether that's with a big provider or with a startup, it's going to be critical to make sure that people component is there, too?

Josh Galvan (8:39):

There's so much to consider. And one thing that I think you're suggesting without saying directly, is that AI is here to stay? This isn't a technology innovation du jour and something companies and individuals are just tinkering with; it really is the latest ultimate combination of existing digital or emerging technology building blocks. That so far, we haven't fully been able to capitalize on. And now, with AI, we can do that. You talked about large language models, about automation and some other things. And I'm assuming you agree that it's here to stay. And we're not just playing around. We're coming up with real business solutions that are going to stick. So that's exciting.

Jason, I what to switch gears a little bit. We get to talk about something that maybe is a little closer to my area of practice on a day-to-day basis. And what I think is a very interesting part of the AI discussion. As with any other technology innovation, there's always a creeping suspicion of risk. And in the wide array of risks out there there's data, confidentiality, privacy, cyber security, theft and so on. And in your area as a professional with many years of experience in cyber security what should we expect from the kind of the core strategies companies are implementing around how to protect AI, how can companies build counter measures to go against these risks that present themselves, some of which we know about, some of which are still manifesting as new types of risks as a result of the new technology innovations we're seeing from AI any thoughts on that, Jason.

Jason Haward-Grau (10:20)

Josh, thanks very much. Great to be here. It's fun to be on the security end of AI right now. I think it's important to start with the preamble, if you're okay with that, just kind of taking it back a step because one of the things that cyber security professionals are used to is being part of the conversation and the dialogue. As things are developing, so technologies develop as we are growing our strategies as businesses. And Anu talked about this a lot around the disruptive

potential that AI has and what we're starting with to a degree is industry 4.0 on steroids. If you think about it and follow the analogy, I do this with a lot of clients from a conversational standpoint, if you think industry 1.0 was very much hands to machine, it was the aggregate, the agrarian farming approach. Then we moved into 2.0, which was the initial infrastructure, technological, the development of railways and telegraphs and the structure of communications. Then 3.0 came along as a digital revolution and that was very much in the 1970's. And now 4.0 on steroids is basically being driven through the opportunity that AI gives us, why is that important? It's because of the velocity of change. One of the things that cyber security professionals are starting to see is how do we secure something that has such a rapid response to it and what we're starting with and it's something that's a new kind of tide. If you break it all the way down as it is incredibly complex and clever and agile as it is, it's effectively the development of software. It's the development of software and the collaboration of knowledge. And what we're tending to see right now is two key areas where security is getting into the conversation.

Jason Haward-Grau (11:59)

Ensuring that we're able to have a conversation around, what do we do with this technology and how do we do it. So, when we start thinking about how we secure it, it's about making sure that the software development life cycle that incorporates AI is correct. It's ensuring that we don't have a solar wind which are inherently part of the equation, or we reduce the risk of them happening, because it's always the potential. And Anu also touched on this that there's a range of different players. And it's incredibly evolving with 700 or so new organizations exist and have suddenly sprung up in the last couple of last couple of years. Understanding the diligence of your supply chain if you think about it, you have a large company, maybe an oil and gas coming, maybe a chemical company, maybe a power utility, maybe a luxury fashion organization. If they're reliant on a startup to build their AI model that they're going to use for logistics, operational integration, and planning with retail for example, you need to understand that organization and you need to understand how that security is going to evolve. So, pushing the security down into the supply chain is crucially important. And let's not forget also, that nation states are really concerned about how rapidly AI is evolving and the potential for misuse. So having a good cyber security strategy in place at the get go helps.

Josh Galvan (13:19):

That's fantastic, Jason. And clearly, the innovation implementation operation of AI solutions is challenging, not just technology, but other functions in the business, and that includes external parties and partnerships to evolve, to upskill, to rethink and as it relates cyber security and risk in general. That means, demanding a seat at the table and an integration with the teams developing and implementing and operating those solutions, so that considerations around cyber are on top of things like legal ethics, compliance, enterprise, risk, and third-party risk are all considered right from the start. Thanks for sharing that, Jason, really appreciate your insight, as always. Just quickly turning our attention to one final topic. And it's really the operational aspects of AI solutions. This, I suppose, is where the rubber hits the road, as it were, on producing a business impact or business outcome, or changing the way in which businesses are conducted, and doing that through AI solutions. Anu. If you don't mind, can you talk a little bit about how you are seeing companies are or will be working with their partners to run and care and feed for AI solutions over

time. I will ask this question this way, because over the last 8 or 10 years we can think about a lot of emerging technology digital solutions, lot of which were rather democratized, and citizen developed and lightweight. Some of them having short shelf lives. So, I'm curious about that aspect, among others, as it relates the operations of AI solutions, running them, replacing them, enhancing them, getting rid of them. What is that life cycle looking like, can we even anticipate that yet?

Anu Puvvada (15:10):

There's definitely many considerations here. Josh, I just remember 5 years ago, when I was working on implementing AI, training was just very costly. It's very people based almost cost prohibitive. The technology just wasn't advanced enough at that time to be able to just take it across entire workflows. And you know, generative AI has been in development for the last decade. And what's really changed is that it's easier to get into the hands of employees. Now, you have this already trained model that you can leverage almost instantaneously. Definitely stage one is around implementing generative AI in the company. And it's important to know that it's a major digital transformation project and it breaks down across people process and the technology itself, and all the implications that come with those and the operational aspects that come with that but when I think about it, I think of stage one is just implementing already existing large language models and using them in certain use cases and letting people be able to use it to draft emails and basic things that they do every day. Stage two is around taking the differentiated data and knowledge of a company and basically incorporating that into the artificial intelligence. And that requires a company to really have their data in order, understanding the pipelines of data that they have, what is the data that they have and how can they leverage it to create differentiated outcomes in the market. You also have to think about the people side of it. And do you have the right change management in place? I see a lot of these type of technology initiatives as you noted Josh, they died on the wine right because people are not adopting them. So, you really have to think about that. How do you train your people to optimize and draw the best answers from this model because it's very conversational. So, it's important that people are asking the right questions, asking them the right way, being trained on how to do things. Like my team right now, we're using something called super prompting in order to draw the best knowledge out of these models. It's really an investigative questioning technique and many times that is not incorporated into training for many companies. And that's one of the main things that as they interact with these models, it's very important to know but all of this technology is really transforming the way that we work. And one of the things that's been top of mind for me and what I've been thinking about is a new form of intellectual property has emerged from this and that form is really how people think and apply that thinking in the context of their organizational process and culture. And that's where differentiated knowledge based intellectual property lives. And so, companies need to be thinking about, how do they capture and protect that knowledge. And it's different than knowledge that sits in a PowerPoint presentation. It's the how of the knowledge. And that's emerging out of this and it's important to capture prompts and be able to harness the best ones so that you can take your company to the next level. There's obviously the responsible AI part of it, ensuring that you have the right governance and guardrails, making sure that you're mitigating hallucination, bias in response, threat actors on the models. And I talked about it earlier, training people. It's easy to just rely on the answer that comes out of a large language model but it's important to remember that it's a prediction engine. It's not thinking it's not reasoning. The humans do that piece of it and so we have to train humans in the loop to ensure that they are actually doing that piece of being critical about the information that's coming out. So, it's a very holistic digital transformation project. And you know, change management is really at the core of it.

Josh Galvan (19:34):

There's so much to unpack there, and I know we're limited on time. But I really appreciate the highlights. Jason, Anu talked about a few things that I know are right in your area of the world cyber information protection. We talked about a number of things there, any views from you, as it relates the cyber operations.

Jason Haward-Grau (19:56)

Yeah, I mean, we're already starting to see that a technology can be used for good or ill. And we're already starting to see from a cyber response and recovery perspective, misuse of AI. We're seeing AI based attacks. We know that bad actors are learning how to apply AI models to their own engineer with their own ingenuity and research, to actually increase cyberattacks. So that's one aspect that's interesting. Anu, you touched on something which I think is really important. It's the data that is becoming more relevant to how we operationalize AI, the securing of that data, the confidentiality risk around that data, the lack of governance around AI. It's not a super cyber security response capability, it's not something that cyber operations per se worries about, but it's something that cyber professionals are really concerned about. Because if you think about it the last survey that I touched on was 11-12% of data inputted into Chat GPT is confidential. I had a client who contacted me recently who wanted to understand what strategy they should be putting in place to simply limit people's ability to use Chat GPT because they're asking questions around intellectual property and critical engineering schematics which were being uploaded into Chat GPT. So, there are small instance of data leakage without that kind of coherent strategy that needs to be addressed. And I think the other part of it, which we haven't got time to touch on it today, but something is interesting when you see government actively reacting to it rapidly, the EU AI act is coming out. There are about 10 or 11 different legislations on the books or in process to provide framework and structure around AI, because it's so rapidly evolving. And, Josh, you touched on this at the beginning, cyber needs to be at the beginning of the conversation. It's not just the heuristic models and the implications for the organization from a how do we transform and change. There are also real-life implications around things that we couldn't do 3 or 4 years ago that we could potentially use AI to do today. Integration between autonomous vehicles and local municipalities is something that's already being done. The question of whether or not that the data that you're gathering, maintaining and assuring as part of that integration the models that are being used wasn't considered when the design was put together. So, there was an immediate potential leakage of, would you really want to know if you're driving from Houston to Galveston, along I- 45, that all the municipalities that you're driving through also happen to know that you're driving through them? And would you be confident that they are able to take appropriate measures to secure the information that relates to you and your vehicle, because the funny thing is, vehicles are no longer cars on their own, they are entertainment platforms, and they are work platforms. It's a very dramatic, rapid shift. So, we're definitely having to think a little bit more holistically, and shift left in the process. As I said, to ensure that cyber security is thought of at the get go, just because you can do something doesn't necessarily mean that you should do something. And if you think about the sheer scale of this, I was talking to two of the large infrastructure and data center providers recently. They are being tasked to provide Colo type data centers closer to decision making processes for organizations because latency is an issue in cloud. They are seeing the great journey that we had over the last let's call it 10-11 years to migrate to cloud and get rid of our own data centers and operate everything in individual clouds is actually not necessarily as effective as it needs to be. How do you secure that? So, we're seeing across a range of operations the impacts and the challenge

is always the same. If you can think about it, plan for it, have an approach toward it you are far more likely to be successful, because, if you think about it, strategy leads into design, design leads into consideration of data, architecture, and infrastructure, thinking about how you're going to model, what is this thing going to do? How is it going to do it? How do you test it? How do you ensure its secure? Anu, you touched on this a lot. There's an inherent potential risk of bias. We're seeing some of those activities already in some of the public AI solutions that are out there where bad actors are looking to try and skew the models in their own direction. She made a great comment which 100% agree with, which is the reasoning and rational component of any AI is not the AI, it's the human being. But if you're relying on it to provide you with material for decision making, therein lies the challenge. Does that make sense?

Josh Galvan (24:35):

That's great, Jason. Thank you both Anu and Jason. I really appreciate your time today. I'm sure our listeners have also benefited tremendously from your insights. There's so much to talk about here. We could go on for hours. I won't be surprised if your email boxes and cell phones are already rattling. I really hope you have a chatbot, or some other mechanism that's AI enabled helping you manage it all. I know you provide some meaningful insight, support, and strategic thinking as well as a technical assistance to companies in their own journeys in AI. This wraps up today's Houston Trend Talks episode. I really appreciate everybody tuning in. I hope you'll listen to another edition of Houston Trend Talk coming off the wire soon. You can connect with us on LinkedIn and other online KPMG resources. We look forward to seeing you next time. Thank you so much. Bye, for now.

Questions? Contact:

Joshua Galvan
Advisory Principal
T: 713-319-2082
E: jgalvan@kpmg.com

Anu Puvvada
Managing Director
T: 713-319-2000
E: apuvvada@kpmg.com

Jason Haward-Grau
Advisory Principal
T: 713-319-2079
E: jhawardgrau@kpmg.com