Architecting the Cloud, part of the On Cloud Podcast

Mike Kavis, Managing Director, Deloitte Consulting LLP

Title: **AI: what it is, isn’t, and how to succeed at it**

Description: Artificial intelligence is all the rage in nearly every sector, but many people still can’t adequately define it, especially vis-à-vis its complement, machine learning. There are also concerns about algorithm bias and failure rates with AI projects. In this episode, Mike Kavis sits down with rp2ai Research founder and CEO Kash Kompella to discuss his book, “*Practical Artificial Intelligence Enterprise Playbook*.” The pair discuss AI—what it is, what it isn’t, bias and ethics concerns, and how to get (and keep) AI projects on track. Kash’s take is that humans need to be constantly in the loop and that governance plays a critical role in preventing both bias and AI project failure.

Duration: **00:29:22**

Operator
This podcast is produced by Deloitte. The views and opinions expressed by podcast speakers and guests are solely their own and do not reflect the opinions of Deloitte. This podcast provides general information only and is not intended to constitute advice or services of any kind. For additional information about Deloitte, go to Deloitte.com/about. Welcome to Architecting the Cloud, part of the On Cloud Podcast, where we get real about Cloud Technology what works, what doesn't and why. Now here is your host Mike Kavis.

Mike Kavis:
Hey, everyone, and welcome back to the Architecting the Cloud Podcast, where we get real about cloud technology. We discuss all the hot topics around cloud computing, but most importantly with the people in the field who do the work. I'm your host Mike Kavis, chief cloud architect over at Deloitte. Today I am joined by Kash Kompella, CEO of rp2ai Research, a global technology industry firm. So, Kash, welcome to the show. Today we're going to talk about a book you wrote called, “*Practical Artificial Intelligence Enterprise Playbook*.” I pulled a few areas out there that interest me and we'll discuss that. But tell us a little bit about your background.

Kashyap Kompella:
Thanks, Mike. Happy to be here. So, as you mentioned I run an industry analyst firm called rpa2ai. So, as industry analysts we analyze emerging technologies, with particular emphasis on artificial intelligence and automation, what's their impact on different industries and what's their impact on

different functions such as HR, marketing, finance, operations, et cetera. We do a lot of interesting work, but I will highlight a couple of things. So, we do a lot of work with private equity companies. We support them when they're investing, or looking to invest in AI companies, or we assist corporate M&A teams when they're looking to acquire AI companies, so that's one of the things that we do. Another thing that we do is we conduct a lot of master classes for managers on the essentials of AI and how they can use AI in their businesses, including AI ethics and AI governance. So, our network spans India, US, and UK, and a large part of our work actually has been in the US.

Personally, my experience is not just in advising companies and businesses in AI, but I have founded and burned myself in doing an AI startup. I've run P&Ls and been an advisor to CTOs, hands-on data scientist myself, et cetera. So, that's my background –

Mike Kavis:
You've got the scars to prove it, right? *[Laughter]*

Kashyap Kompella:
Yes, yes. Unfortunately, so, yeah.

Mike Kavis:
So, one of the quotes from your book I think is going to play into what you talk about – so you analyze a lot of the AI companies for investors, and one of the things you said in your book, and I see the same thing, is, "The first myth to bust is most AI is not really AI at all. Most AI-labeled products are machine learning at heart." So, discuss that.

Kashyap Kompella:
Hmm, that's a very interesting question. So, I guess we are making a distinction between artificial intelligence and machine learning. So, the term artificial intelligence, it was coined in 1950. Since then, its main goal has been to develop machines that have intelligence that's similar to that of humans. But machine learning, it's a branch of artificial intelligence, but it works slightly differently. Machine learning is learning by example, and data is central to machine learning. And really, the data set that is used to train the machine learning model, it determines the predictions, it determines the outcomes it produces, et cetera.

So, the training data set is really the machine learning model's reality. Using machine learning, for example, we're able to build systems that perform extremely well on specific tasks, but they still have strong limitations, so they're data hungry. And at the same time, if there is a deviation, they're unable to make sense of situations that are different from their training data, or examples that they haven't seen before. And – but artificial intelligence, when people think of intelligence, that's more about, can you create an AI that is capable of solving tasks that it has never seen before without significant involvement from human researchers.

So, to sum up, in a sense – there is a quote that I'd like to read out. It says, "Machines might do some things as well as we do them, or perhaps even better. But they would inevitably fail in others, which they would reveal that they were acting not through understanding." So, I think this quote sums up the state of machine learning, the systems that we're producing right now. But this is actually a quote from the 1600s from the French philosopher René Descartes. So, what's beautiful about this quote is that it brings up the flexibility, the adaptability, the generality that we associate with intelligence, and that's not machine learning.

Mike Kavis:
Yeah, the quote that's stood the test of time, huh?

Kashyap Kompella:
[Laughter] True.

Mike Kavis:
Yeah. So, one of the other challenges I see in this, and you have a whole chapter dedicated to this called "The Dark Side of AI," is kind of the bias, the blind spots, the ethics of all this. So, what's your take on – you obviously wrote a chapter on this and you have a lot of thoughts on it, but what's your take on the dangers or the dark side of AI?

Kashyap Kompella:
That's a fascinating question. I actually write a column called "AI Ethicist" for Information Today (infotoday.com), and I think about this topic a lot outside of whatever short treatment to the topic we have given in our book itself. So, if you think about it, in the news you must have read a lot of examples or instances of AI having gone wrong. So, the primary problem is that AI systems are currently being created by technical and business teams that are looking to solve a narrowly-defined set of questions. And these teams, and the systems they create, don't know what they don't know, so that's a problem. For example, if you take the case of facial-recognition systems, they make mistakes, and the percentage of mistakes is higher for certain demographics groups such as women and minorities.

And that's usually because the data that is used to train those systems, it's not balanced. It's not representative. But when you ask a common person or a regular user, in the users' minds, AI systems are infallible. So, what ends up happening is AI systems are actually perpetuating and accentuating the existing inequities in our world. And what's even worse is that such biases and errors are automated, and sometimes there is no recourse or any redress mechanism, reversal mechanism. So, that explains a lot of unease about the widespread deployment of such AI systems without the guardrails and the safeguards and controls.

And it's not an easy thing to solve, because ethical questions have no easy answers, because unlike an engineering problem, there is no universal objective function that you can optimize. There is no singular objective function. There aren't any universal decision rules that you can just apply. So, what I'm saying is you cannot simply code your way out of this thing. So, that's the challenge that is becoming increasingly important as we go forward, as AI becomes more widely adopted.

Mike Kavis:
So, does the risk, the ethical risk and the blind spot risk, does that prevent a lot of companies from using these types of things externally and using it more internally? So, for example, I don't remember which company it was, but some company put some kind of bot out there on Twitter a couple years ago and it turned – it started saying some really bad things, right? It started picking up patterns and saying bad things and they shut it off. Is there a tendency, because of the ethical risk for companies, to use these technologies more internally than outward-facing to their customers?

Kashyap Kompella:
Not really. I think the biggest risk there is there isn't enough exposure or awareness of such risks, so that's a problem. Say if a chatbot goes out of whack and starts spewing offensive language. It's not really offending or creating real-world harm, because I mean, that's an everyday (Inaudible), right? But there are certain examples, where there can be certain real-world consequences. So, one of the classical examples of thought experiments is what's known as a trolley problem. So, the setup is something like this. Imagine an out-of-control trolley car that's coming down a track towards you or towards say a certain number of people, five people. They're chained up and they cannot get out of the way. Now you're a bystander. You're out of harm's way and you're standing next to a switch that can change the tracks for this trolley. But on the other side of the track when you switch that, there is a person who is similarly chained. So, there are five people who are chained and there is one person on the other track. So, what do you do? Do you do nothing and let the trolley crush those five people? Or do you pull the lever and cause harm to the other person? So, what's the right thing to do?

So, traditionally how – I mean, this is an example in machine ethics, and traditionally how we have tended to approach is that – this is more like different schools of philosophy. There is utilitarianism that says that actions are right if they benefit the majority of the people, or another thing called deontological ethics, in which the rightness of actions is not linked to such cost-benefit analysis. So, actions are either good or bad according to a set of rules regardless of consequences.

So, what's good really varies based on your outlook, your philosophy, et cetera. So, this trolley problem looks like an interesting thought experiment. But in effect, there are a lot of trolley problems that are already happening right now. For example, who receives organ donations? Who receives a parole in the criminal justice system or who doesn't? Whose loan gets approved or not? Who gets hired? Who doesn't get hired, et cetera? So, if our assumptions are that – so that the programmers, the programmers and the developers who are creating these systems are not making these choices. They're optimizing for accuracy of prediction based on the data that is being submitted. So, the way out of this is really – as they call it as like we have to make these systems AI – the AI systems should be ethical inside, but how do we make that? We need to bring in that governance function, that quality assurance function which can confers these choices. And that suggests that the teams that are creating these solutions need to be diverse. They need to include not just diversity from a demographic point of view, but diversity from a worldview point of view.

There are a lot of examples. For example, in the trolley problem, you can have a very different trolley problem in which there is a senior citizen versus a child; you have to make a choice between them. It also depends on culture. Some cultures have a lot of respect for senior citizens and they say that that should be given priority, while some others, they'll say, hey, it should be the young who have a lot more years ahead of them who should be given priority, et cetera.

So, these are tough choices, and I don't think there are any easy answers. Yes, we do have some emerging approaches in terms of can we fix the dataset that we're using. Can we fix or look at the harms caused by the algorithms that we are producing, et cetera? I know I'm sort of painting a very broad picture, and that is representative of the problem we have at hand. Is this problem very unique to the AI field itself? I mean, probably yes to an extent because of the pedestal in which common people place AI because of representations in the media, media narratives, and because AI means different things to different people.

Mike Kavis:
Yeah, the trolley example is pretty interesting because there's no good solution, right? If you don't do anything, it's a disaster. If you switch it one way, it's a disaster. If you switch it the other way, it's a disaster. And it comes down to people's interpretation of what was ethical or right. That kind of puts this in a context I never thought about. In the past humans made those decisions and now it's up to machines, and it's probably a lot easier to blame the machine, right? So, it's an interesting dilemma there.

Kashyap Kompella:
Yeah, it's an interesting dilemma because – this comes up, I mean, quite often. People love to discuss this in the context of (Inaudible), so something like this, what will happen? Where does the liability lie?

Mike Kavis:
Well, there's another part of it. So, if we switch it to a happier use case, let's say one that's making one of three business decisions, the one thing that might be missing is a human experience. So, you may have an experienced leader who, even though the data may favor one path, they just know from experience or from knowing that customer that path two might be better. Does machine learning, AI figure that out? Or, just the whole dilemma here is should we do it fully automated, or should it be augmented where the machines are telling us and then we make a choice? And I think it really depends on the use case. I was just wondering what you were thinking there.

Kashyap Kompella:
That's correct . So for certain decisions, see, do you have skin in the game or not? So, you mentioned an experienced leader. I think people who have skin in the game, they tend to make certain kinds of decisions. I mean, that skin in the game could be in terms of reputation, in terms of monetary incentive, or in terms of the impact they want to create, et cetera. So, I want to take a step back and talk in terms of what does it mean that machines and algorithms are achieving human level of performance or exceeding human level of performance in certain tasks, say making a decision. That means that they are right 80 percent of the time, that they're right in some cases 90 percent of the time, et cetera. So, the tricky issue is really which is the 80 percent of the happy scenario and which is the exception? So, I think in all machine learning workflows it's programmed to include exception handling, or human in the loop and have a reversal mechanism if things go wrong.

So, there is an example – I don't know if you followed it. Because of COVID certain schools in the UK were not connecting their exams, their exams for graduating students, I think 12 high school students. But next year's college admissions are not going to stop, and they were going to be given their marks and grades. So, the board has relied on an algorithm to give grades to the students, to give them how they're going to perform, et cetera, et cetera. And one of the factors that was included in the algorithm – this came out later – was that school which the students were attending, because typically there was a correlation that said that students from certain types of schools, which tended to be lower income schools, are not going to have higher grades.

So, this is a pattern that the machine learning model has figured based on the data, et cetera. That obviously put a lot of students at disadvantage, even high-performing students from such schools, et cetera. So, what happened was they did not factor the people who were using the systems who were implementing these systems, they did not factor a human in the loop in an exception process. If a student had a genuine issue that needed to be addressed, there was no way they could do that. I mean, they went back and fixed it later, so it's not just in the design and implementation of systems, but how do you actually use them as a user? How do you actually build in, like I mentioned, those guardrails? That becomes important, even in the business context. Is this the right decision or is it going to lead to outcomes that we haven't expected?

Mike Kavis:
Yeah, I'll use a recent baseball example. So, there was just a World Series. I live in Tampa Bay and the Tampa Bay Rays were in the series. And it comes down to game six. They've got their star pitcher on the mound. And their manager is up for manager of the year because of his use of analytics and how he managed his pitching staff. So, he has all this data and he just makes decisions based on the data. And the one datapoint for this pitcher is the third time around in the batting order that the teams tend to hit him. So, he's pitching a one-hitter in the sixth, right? And a guy gets a hit, and he uses his analytics and takes him out. Now this is the World Series, right? And the next three guys coming up have struck out almost every time up, so this guy is on his A game. But he's trusting his analytics, because that's what got him there, right? The next three guys get hits, and he went from being manager of the year, now the whole city's piling on him, right? So, that's an easier example to consume, but he got there through 60, 70 games of managing based on analytics, but it hurt him when it really mattered at the end. So, just an example that hopefully sports fans can relate to.

But I want to pivot, because a lot of what your book is about is there's a lot of failed AI projects and you're kind of putting a playbook or guidance on how do we prevent these failures. So, first, what are the common reasons for AI projects failing? And then how can we make them more successful?

Kashyap Kompella:
Yeah, I mean, that's a great question. I feel that there is a lot to learn from observing failure just as much as observing best practices of success. But I want to take a step back before addressing AI failure itself. I mentioned my startup, and 90 percent of startups fail. If you talk about mergers and acquisitions, about 70 to 90 percent of acquisitions fail. If you look at large-scale digital transformation efforts, a huge portion of them, majority of them do not succeed as expected. So, in that sense, failure of AI projects is not very unique. I don't have exact numbers, but my sense is up to 80 percent of the AI projects fail. And it's interesting to think about how does this failure manifest? Say it manifests as failing to make a meaningful impact either on the project objectives or bottom line or top line, or there is a lot of disappointment because these projects lie abandoned, unable to be scaled.

So, when I look around and see what are the common factors for this failure, for this lack of clear success, they fall into five broad buckets. One is business, data, people, process, and technology. So, let me take each one by turn.

Let's start with the business. So, there are very unrealistic expectations about AI. I mentioned about media narratives of AI being portrayed as a superpower or superhuman. "Anything that humans can do, AI can do better," is the prevailing narrative. So, because of that, there is a lot of pressure from top management to say, "Okay, let's implement AI and the thing will be fixed," but unfortunately that's not how it works. So, being able to identify those use cases, having realistic expectations in terms of what should be the project's scope, what should be the project's tangible benefits, et cetera, selecting the right use cases, that becomes very important.

So, the next dimension is around the data. So, I get a lot of calls from companies saying, "Hey, we need to have an AI roadmap for the next year," which usually happens in the last quarter. *[Laughter]*

Mike Kavis:
Yeah, use that budget, right?

Kashyap Kompella:
Yes, yes. So, we're like, "Okay, that's fine. Let's do it. This is a great thing." But as you start digging into what the company wants to achieve, do they have their ducks lined up, et cetera, you find that they do not have the data that is required. I mentioned machine learning. I mean, its reality is the data set that you have. The data is either of low quality, it doesn't exist, or it is in different systems that do not talk to each other. Data may not have been collected, or even if it was collected it's not in a form that's available to machine learning training, et cetera. So, as we're saying, "Do you have this, do you have this, et cetera?" More often it comes down to the fact that a lot of companies, I mean unless you're a Silicon Valley tech powerhouse that has the chops and the data, which is more readily available in a digital format, companies are better off putting in the foundational technical pieces in place for data. So, that's the next obstacle that needs to be cleared before AI projects that can be successful.

Then, I mean, all of us know that the skills, the people dimension, skills are in short supply. So, there is some statistics that say there are three jobs for every machine learning engineer out there. The situation has changed slightly, not too much, with COVID, but still the demand is greater than the supply. But that is slightly changing. There is a lot of interest and awareness and people training themselves, getting certified, et cetera, so that's changing. But people constraint is going to remain for a little while, and in fact I think the biggest opportunity in AI is going to be around consulting and in services.

Another aspect of the people dimension is also the governance aspect in terms of approving the pipeline of business cases that need to be produced. What is the staffing that is required? What should be done in house and what should be done externally through consultants, et cetera? So, we found that here, having a strong central group that oversees all of these functions, standardizes the machine learning projects that can be taken up, et cetera, that's really useful, starting with the people who take stories of success from one group to the other. What has worked? What hasn't worked? What should be the building blocks in place before you become ambitious, et cetera? That's going to be useful.

Then there is process. Process is not unique to AI itself, but there are some subtle differences. I mean, you don't want to handle AI projects as your plain, vanilla IT projects, because in an IT project, you're sort of done at some stage. You update to a new system; you're done. Obviously there is continuous integration, continuous development, but they tend to be incremental rather than too much that is disruptive. Like you implement an ERP system and you expect it to be working, you leave it as is, but in an AI system after you deploy it, you constantly need to tend to it. As the real world production data keeps coming in, the model's performance may degrade. So, your job is never done really, so you need to have different models in place, different monitoring mechanisms, measurement mechanisms in place from a process. And the important question is what methodologies do you use, what frameworks do you use, et cetera? Those are all on the process side.

On the technology side, that's the last dimension I wanted to mention. There are two aspects. The first is the capabilities of the machine learning platforms themselves, and the second thing is some sort of tuning and instrumentation to move from the pilots to the production. So, this is the realm of Machine Learning Ops or some sort of a DevOps for machine learning. How do you take your model that your data science team has produced and move it to production and bring it back, integrate it with your entire lifecycle of technology projects, et cetera?

The good news here is things are improving, looking up. This is a very happening field in this AI domain. If you look at the examples of some of the Silicon Valley companies, all of them really have had to build their own MLOps. Facebook to Ubers to Twitters. Most of them have released these frameworks for people. But now you have the average enterprise for a regular bank or a telco, et cetera, these are available – products are becoming available for Machine Learning Ops. And, I mean, technology selection platform – I know you're a cloud expert. So, most of the popular clouds, they have machine learning as an API or machine learning as a service, but they tend to be broad and not as much wide. So, you may have to sort of figure out whether the technology platform really suits your requirements and use case, or is it still not ready yet?

So, I'll stop there. So, those are some of the common factors that we've observed and what are the things that you need to consider to increase your odds of success with machine learning projects.

Mike Kavis:
Yeah, what's interesting is everything you said is just the same thing I say about adopting cloud, right? It's the same thing. There's people, process, technology, and the same challenges. And one of the things I've written about recently is that I started in mainframe and there wasn't a lot of change till we got to client server, and then a few years later there's thing called the Internet. You know, there was a lot of time between the big transformations. Now it's like every six months there's something transformational coming, and one of the problems I see is even the companies who see these things as a transformation and invest in the people, process, tech, they're doing it in technology silos, right? So, there are people working on an operating model, a new way of work for machine learning, but they're not talking to the people who are doing the same for cloud, and they're not talking to the blockchain team. So, now it's just like we're changing in silos and just making things ugly, and I don't know the answer. Part of me says there needs to be a center of change, but boy, that sounds like a lot of bureaucracy.

But the point here is that we're absorbing change from all directions like never before, and change is hard. So, we buy all these tools and technologies but it's really hard to put them in play and actually create business value of them. I don't know if you're seeing that, a lot of places where clients are just trying to implement too many transformational things at once and getting none of them to the finish line.

Kashyap Kompella:
I think I agree with you that change is the only constant sort of theme in the world of technology, and the pace of change has accelerated. I think what also has changed is there is this new fear of missing out. So, that hasn't been the theme in the previous waves of technology, I think. So, I mean, my thesis, and I have seen this play out, is this transition to a newer, AI-first or a more widely-spread AI is going to happen over a couple of decades. It's not going to happen overnight. I've been an industry analyst for more than ten years now, and I see that the kind of things, the trends that succeed, that people keep talking about, they actually happen after five years. So, ten years ago cloud was all the rage, but I haven't seen a lot of spend on cloud at that point in time. So, when I look at the CIO priorities this year, the last year, or in the last three years, 80 percent of that has been going to cloud. And all the action and the real spend is in things like more foundational things like cloud, but the talk is on newer, shinier objects, like blockchain maybe, or AI.

So, one example is that if you look at the financial results of a company like Amazon AWS bulk comes from core cloud, while the revenues from the machine learning and AI services is probably much, much smaller compared. So, the challenge for leadership is to sort of keep an eye on the future, but at the same time don't forget the core basics, et cetera. I don't know if that answers your question, but that's balancing the short-term with the long-term aspect.

Mike Kavis:
Well, I think becoming really good at cloud for a company and really good at embracing DevOps concepts, I think those things are stepping-stones for making these other technologies and transformations successful, like AI, machine learning. I think you don't need cloud to do this stuff but obviously you can really accelerate if you leverage cloud services. But, there's a lot of change between cloud and DevOps to get to a mindset of continuously learning, higher degrees of collaboration, being accepting to change. And I think if companies are three, four, five years into that journey and doing well, I think embracing you know, whether it's edge computing or AI, it becomes easier. But the problem I see is a lot of companies are trying to tackle it all at once. They're starting their cloud journey at the same time they're starting their AI journey at the same time – and it's just madness, so –

Kashyap Kompella:
I think there are two things, aspects – I mean, if I just want to latch on to what you said. That is, I think previously technology was not seen as a core competence of a company, but now that is changing, and a company needs to have a lot more technical chops through the kind of things that you mentioned, having that agile mindset, being able to embrace newer things like cloud. And at the same time, they also need to be cautious about not accumulating a lot of technical debt, because you won't be able to move as fast when it's needed if you have a lot of things broken.

Mike Kavis:
Yeah, I agree with that. This could be a whole other podcast, this topic. Unfortunately, we're kind of out of time. But I really appreciate you coming on today and sharing your insights on AI and just the future of IT as we kind of wrapped up here. Where can we find you? Do you have a Twitter handle or – I'm sure you've got a lot of good content out there. Where can we find all this?

Kashyap Kompella:

Yes, I think the best place to find me is on LinkedIn, Kashyap Kompella. I have a lot of content out there, not just on AI but a lot of other topics for business and technology leaders, so please find me on LinkedIn.

Mike Kavis:
Will do. Make sure you check him out there. I looked at a lot of that content, some good stuff. So, that's it for our episode today of Architecting the Cloud. To learn more about Deloitte or read today's show notes, head over to www.DeloitteCloudPodcast.com. You'll find more podcasts by me and my colleague Dave Linthicum just by searching for Deloitte On Cloud Podcast on iTunes or wherever you get your podcasts. I'm your host Mike Kavis. You can always find me on Twitter, @MadGreek65, or e-mail MKavis@Deloitte.com. Thanks for listening and we'll see you next time on Architecting the Cloud.

Operator:
Thank you for listening to Architecting the Cloud, part of the On Cloud Podcast with Mike Kavis. Connect with Mike on Twitter, LinkedIn and visit the Deloitte On Cloud blog at www.deloitte.com/us/deloitte-on-cloud-blog. Be sure to rate and review the show on your favorite podcast app.

Visit the On Cloud library
www.deloitte.com/us/cloud-podcast