



For Cloud Professionals, part of the On Cloud Podcast

David Linthicum, Managing Director, Chief Cloud Strategy Officer, Deloitte Consulting LLP

How is the cloud accelerating the adoption of AI?

Join host David Linthicum and guest Tom Davenport, a well-known thought leader in analytics and AI, as they discuss the evolution of artificial intelligence and how cloud computing is helping to accelerate the practical application of AI technologies. You'll also hear Davenport's perspective on the most surprising findings from Deloitte's 2018 State of AI Survey, plus highlights from his new book on how enterprises are winning with AI today and what to expect in the near future.

Duration: 0:25:36

David Linthicum:

Hey, guys welcome back to the podcast. We got a special guest today, I mean, a celebrity in the AI space, somebody who I met actually three weeks ago at Deloitte University and he is a great guy and also read his book on the way home and it's a great book and so happy to have him on the podcast, it's Tom Davenport. Tom is the President's Distinguished Professor of Information Technology and Management of Babson College, the Co-Founder of the International Institute for Analytics and a

fellow at the MIT Initiative on the Digital Economy and a Senior Advisor at Deloitte Analytics which is how I got—linked up with us. Do I miss anything? Tom, you are kind of tell us what you are working on now, what have you done in the last month? What's kind of your primary day job?

Tom Davenport:

Well, I'm still writing articles about some AI and the Application of AI, doing a fair amount related to Healthcare and Life Sciences, which I think is a big growing area for this sort of thing. And I don't know yesterday in DC spoke at an analytics and AI conference so that's sort of what I do in general these days.

David Linthicum:

What did you speak about the Analytics in AI conference, what are kind of the cool topics today?

Tom Davenport:

Well, this was a vendor called Thoughtspot that does a couple of interesting things that I think illustrate part of where AI is today. One is sort of search-based analytics so making analytics much easier for people to sort of use and access previous searches and so on, so sort of a Google for analytics and then the other thing they do which they didn't I don't think invent but more and more companies are doing it these days is kind of automatically finding patterns and data rather than relying on humans to do it so it will [inaudible] which identifies a number of insights. The first I'd seen of that was at Salesforce with Einstein Analytics so they acquired a company called BeyondCore that did that but now I hear more and more companies thinking about it. It's quite an interesting idea.

David Linthicum:

So you've been in the IT world for quite a bit, so you've written 18 books and over 100 articles for Harvard Business Review and Sloan Management Review, the Financial Times and many other publications. So what's the kind of changing pattern in AI over the last 10 years, what's kind of new to the market that's really kind of the game changer because reality is, I'm 56 and I did AI back when I first got out of college did LISP and M1 and things like that that and now here we are 2018, and it's kind of making you come back but it's making you come back with different enabling technologies in different platforms. So that's what I've observed but you've been in it probably longer than I have and watched this particular set of enabling technologies, what's changed in the last 10, 20 years?

Tom Davenport:

Well one is certainly the Cloud, which I suspect you'd like to talk about further. Another is these deep learning algorithms that people had sort of given up on neural networks for dead basically and—in part because a couple of MIT guys wrote a book saying that “They didn't want to go much of anywhere as a predictor of a complex phenomenon.” But now I think in certain parts of the world, Silicon Valley, for example, deep learning basically is AI to a lot of people. I don't really view it that way I think it's one appealing technology it turns out that a lot of the things we thought were pretty much dead from the last generation of AI are still around to a substantial degree. Rule-based systems or what we used to call expert systems. But there are a lot of new things, new software, algorithms, new graphics processing units, that churn through the data and basically a lot more data in general. That means we really need AI if we are going to make sense of it all.

David Linthicum:

So one of the things that I hear and this is a good question I get at conferences from time-to-time. It talk about, has really the reemergence of artificial intelligence the derivative such as machine learning and deep learning and things that you talk about in your book, a matter of the Cloud coming along and just making it cheap and better or is this really kind of the technology finding a home that it always we should have had and never really had and basically looking at this as kind of the next generation of cool enabling technologies able to make sense of the petabytes of data we have lying around the enterprises. So which came first, the Cloud or AI?

Tom Davenport:

Well, I think AI but it turned out that the Cloud was really well suited to AI and it only came along a little bit later as the vehicle for doing AI and I don't think we would have had nearly the upsurge in AI that we'd had if it were not for the Cloud but you can do pretty much everything anybody is doing with AI on premise but it's just—obviously it requires more of a financial commitment and not as flexible. And the Cloud has turned out to be an amazing distribution mechanism for algorithms—most

of the leading Cloud providers have made available of set of algorithms, mostly in Open Source but not totally, that make AI much easier to do. It's not that they weren't available but there is certainly much more broadly available because of the Cloud.

David Linthicum:

So it's kind of both, in other words it's kind of the perfect storm of we need this technology now they are trying to get data under control and really trying to discern the patterns of the data and become deeper at analytics and trying to build things that are going to be much more reactive in nature and do some learning models in nature. As well as the fact the Cloud Computing kind of came along and just kind of made it easier to do we can leverage it on demand. I can start up my AI skillsets with just a credit card these days versus back when I first got out of college I mean we were building things that literally cost \$100 million of datacenter space just to get simple questions answered. So, it's a huge change but, do you think that enterprises are leveraging the technology correctly and leveraging Cloud Computing correctly in AI? What advice would you give them?

Tom Davenport:

Well for the most part, I would say yes, it's a great way to experiment with AI to explore a whole variety of different types of capabilities not just software based but there are GPU hardware capabilities in the Cloud so most organizations haven't had a lot of exposure to them. So they can see how much do these GPU's really accelerate my algorithms and so on. I think if there is any downside it's that it's so easy to experiment that many companies don't really ever get beyond that, they do lots of pilots and proofs of concept and prototypes and so on. But they don't move to production deployment and the fact that it's in the Cloud isn't really what's preventing them but it's just so easy to experiment that the ratio of pilots to production deployments I think is getting smaller and smaller. Production deployment means integration with your existing systems and it probably means changing the way your business process operates and reeducating your people a little bit and so that's a lot of effort. Maybe if you had to commit substantial time and resources to an on-premise solution you might be a little bit more dedicated to doing something in production.

David Linthicum:

I found the big problem that I had when I was leveraging AI, and I've – I've kind of in fits and starts throughout my career is that there were typically misapplied and so in other words we had a problem domain or something or an application we are trying to solve our AI was overkill or not a good fit so things like that. So in other words it was more procedural logic to do just a good job versus building a learning model on the backend even if it leveraged data and ultimately I think a lot of the issues that I had with AI and my clients had with AI was the fact that there was a mismatch and misunderstanding in terms of how to leverage it and back in the LISP and M1 days it was – it was over applied lots of times. People were leveraging it with these odd based systems or that weren't complex in nature and couldn't benefit from a learning model. But these days we have lots of instances of applications, massive amounts of information that can benefit from learning models. So how do you kind of fit the problem with the solution, do you think AI is going to be over utilized in many instances or do you think that we're necessarily going to misapply the technology in ways that are going to be unproductive?

Tom Davenport:

Well there I have certainly heard of a fair amount of that not so much AI in general but the specific type of AI, I think it's important to recognize for any organization embarking upon this that there are a lot of different technologies that comprise AI and so you need to know if your particular use case oriented to some form of natural language processing and even within that there is statistical and semantic natural language processing is it oriented to some form of deep learning based image recognition is it oriented to just the use of logic in which casing a rules if it's of relatively simple problem rules probably still are the best—there was an early on, one of the earliest firms to implement IBM Watson was a health insurance company trying to do pre-approvals for member operations and medical treatments and Watson turned out to be overkill for that problem far more complex than they needed – rules worked pretty well actually. But I think knowing—and even something as simple as robotic process automation—can be a very useful aid to a lot of organizations if you want to easily and quickly automate a process. So, it's important to know what the options are and kind of how the types of use cases throughout an organization match up to the technologies that are available.

David Linthicum:

So with the advent of Cloud we have kind of the ability to have a massive data source. It's not unusual to see data warehouses out there, big data systems that have three or four petabytes which is something was inconceivable even to me 15 years ago,

20 years ago. We have AI systems that layer on top of those, how do we basically build AI systems that are able to make sense of the information – is this about, be able to lop of parts of the database, is this about going through the entire database, is this about doing things in real-time or looking at information is this flowing between source and target systems? What seems to be the major pattern out there in terms of how enterprises are employing AI to make sense of the data, learn from the data, and start discerning solutions from the data, answer problems from the data?

Tom Davenport:

Well, the vast majority I would say of AI applications involved some sort of machine learning and a vast amount of machine learning in businesses is supervised machine learning which means you are training a model with some piece of data for which you know the outcome variable and that is something that's been going on for a while but it's clearly happening more and more often and it applies whether your model is something as simple as linear regression all the way to these very complex with deep learning models with thousands of variables and so on, all of it needs training data to learn on. Some obviously—maybe I don't know if it's obvious or not—but the more variables you have the more features in your models the more data you are going to need. So if you are doing a really complex deep learning model to identify potentially cancerous lesions in radiology images you are going to need an awful lot of labeled radiology images; this one is cancer, this one is not cancer, etcetera that the system can learn on.

So not only do you have all of the cost of storing all that data but you also have to label it and be pretty sure that your labeling is correct, I mean there are some technologies, I was talking to someone from China this morning that talks that uses weekly supervised learning. First, I thought he was talking about something done every week, but it's a less strong form of supervision, and there can be some problems in the data and it can still be quite successful. But in general, this labeling issue is a big challenge for organizations and then you do have, as you say to store it, and then, run your model through it when you don't know the outcome you are seeking which some people describe it as a scoring process. So that's not a continuous process generally and if you want continuous learning, you typically would have to go through another episode of training your model. People often think with AI, it learns continuously from new data, but only if you are willing to use that new data to develop a new training data set and a new model from it.

David Linthicum:

So it's difficult to have these things done in real-time so in adding information to the database in essence, so we have to relearn and re-create the model. Would it be able to update the model continuously?

Tom Davenport:

Correct. We may at some point have continuously updated machine learning models, but we don't have anything like that today.

David Linthicum:

So we did a State of the AI survey or State of AI survey, kind of beyond the survey, what were some of the core things that you found that as we looked and talk to people about Artificial Intelligence, you being kind of an academic thought leader, understanding the space in a very detailed way of course and looking at people's perception to what AI is, what are some of the links that you found?

Tom Davenport:

Well these are still early adopters; this is our second year to do this survey. We found a higher proportion of people this time to really understand it, that was a big factor last time in filtering out potential responses because a lot of people didn't really know enough about what their organization were doing with AI to talk about, this time we found a lot more and they're still very bullish overall feeling like they are getting benefits already, the more an organization has worked with AI, the more positive these respondents were. There were a few changes in sort of how likely people see AI as kind of revolutionizing their organizations in the short run. Last year the majority expected—both our company and our entire industry will be radically transformed by AI in only a few years—but the percentages of that did drop somewhat this year which I think it's a bit more realistic. We certainly still see a variety of technologies; some things like deep learning are going up quite rapidly. We saw in part because we didn't ask about it last year, but we thought 50% basically or 49% said that they were doing AI in the Cloud so, that was certainly an interesting development and then more and more people are using software much of a cloud-based Software that has AI sort of baked in. So I mentioned Salesforce Einstein, SAP Leonardo, Oracle has some AI capabilities

increasingly built in with software, so that's probably the easiest way to do AI, is to use a vendors offering, the data is there, the user interface is one that your organization is comfortable with, and all you have to do is pay a few more bucks every month to have those new capabilities.

David Linthicum:

So what were your kind of feelings and if you are going to kind of abstract the report and some of the top three trends that are occurring now that maybe some things that surprised you, you figured out you knew or understood or things we are going to probably be facing in 2019, 2020?

Tom Davenport:

Yes, I'm happy to do that. One was the issue of sort of risk starting to rear its ugly head in AI and a fairly substantial, I forget the exact number, or you can look it up, Deloitte 2018 State of AI Survey, but a number of people suspected that there would be, if not, if it wasn't already Cyber Security Risk. I think 37% said that they saw some ethical risk related to AI so it's still early but the appearance of risk-related concerns is certainly one interesting development. Another is I'd probably put the rise of the Cloud, and the rise of Embedded AI capabilities and software as one of the top three developments. And the third one was a big surprise to me, when we did this last year, we asked what people had in mind as an objective for their AI work and last year as well as this year they put relatively low on the list, I think it was almost last both years, reducing head count through automation. But this year we had a question saying, "Do you intend to automate as many jobs as possible with AI?" and 63% said yes.

So I thought it was mostly a settled kind of topic that we are going to see much more augmentation than large-scale, automation, and people were eventually going to help educate their employees for how to make them more effective partners with AI rather than just sort of automating their jobs completely, but that finding really surprised me and then we had another finding suggesting that the great majority, something like 80%, said that they intended to hire new employees to do AI-related work, I mean the question was a little bit unclear about what kind of work we meant by that as opposed to retraining their existing employees, only 10% said that they were really dedicated to keeping existing employees and retraining them. So, if you are a human worker, I suspect most of the people who are listening to this podcast are, it's got both of those numbers have got to give you pause, I would say.

David Linthicum:

Let's talk a bit about your book, "The AI Advantage: How to Put the Artificial Intelligence Revolution to Work (Management on the Cutting Edge)." When I read the book, I thought it was very pragmatic I thought it a lot of AI books I've read in the past and some of the manuals, I read in the past were kind of at a such a high level, abstract level, I really couldn't figure out how to make it work in real life, but you seem to cross that boundary in the book. So tell us how you came about it and what were your objectives with the book and tell people what's in it.

Tom Davenport:

Well, this is, I usually write books about how new technologies affect enterprises, a couple of years ago I wrote a book before there was a lot of Enterprise AI on how AI affects jobs and skills and found a number of examples of this augmentation environment that I described earlier, but now there are lots of companies who are using AI within their walls and so the book is about how do you deal with Enterprise AI? What technologies are people using, what kinds of process changes are they trying to create with their AI use cases, what are some of the job-related issues? What are some of the ethical and social related issues involving Enterprise AI, so it's something to be a kind of a broad look at what's really happening now and what might happen in relatively near future with Enterprise AI.

David Linthicum:

One of the things I liked about the book is the first chapter, "Artificial Intelligence comes of age slowly." And I thought that was kind of a very pragmatic approach. You are writing a book on AI but I think being very realistic about how people are adopting out there, so what's your take on that?

Tom Davenport:

Well I tried to contrast in the first chapter this sort of moon shot idea that AI is capable of these incredible transformational feats, of and things like curing cancer or recommending the best investments for financial services customer, or something

along those lines with the more sort of low-hanging fruit types of applications of AI and one of my case studies is Amazon because there is not much doubt that Amazon has about as many capabilities as anybody in the world with regard to AI, but even they find that the vast bulk of their work in Machine Learning and AI is—to use Jeff Bezos’s term—“quietly but meaningfully improving core operations.” So I do think that it’s a revolutionary technology and the word revolution is in the subtitle of the book, but its revolutionary over the long run and clearly evolutionary today and some people may want a more pipe-filled treatment of this subject but I was quite focused on what was really feasible and what’s actually happening today.

David Linthicum:

I liked it. Like a lot of my reviewers say, “David Linthicum is designated buzzkill,” and I think people have to actually put that into perspective. The reality is that technologies are going to have limitations, it’s going to have things that they can do and things that it can’t do and I think it’s good to be honest about what it can’t do, because once we know what the limitations are then we could start working on and basically improving and bypassing those limitations. So where can we find you on the web?

Tom Davenport:

www.tomdavenport.com Not too tough.

David Linthicum:

Okay and you, we can pick up your book in Amazon and any other places books are sold?

Tom Davenport:

I think it’s pretty much everywhere, published by MIT Press.

David Linthicum:

Please pick up a copy of his book I think it’s probably something if you are looking to figure out what AI is, this is the book to actually read. It’s not necessary a primer but it gets in to the details quickly and gives you a pragmatic look at what AI technology can bring to you, and how it’s working in the Cloud, things like that. So, anyway, I want to thank you for being on the podcast and we’ll talk to you guys again real soon, take care.

Operator:

Thank you for listening to On Cloud for Cloud Professionals with David Linthicum. Connect with David on Twitter and 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

About Deloitte

As used in this podcast, “Deloitte” means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.

Copyright © 2018 Deloitte Development LLC. All rights reserved.