Paul Thies: (silence)

Emerging technologies such as artificial intelligence and predictive analytics are dominating today's business headlines and are ushering in a multitude of strategic changes to a variety of industries.

Organizations of all sizes are directing massive resources of time, money, and intellectual capital into AI development, [00:00:30] in a bid to realize market leading leaps forward and technological capabilities.

But where is the zeitgeist that is artificial intelligence development headed? Welcome to If/When Jacobs series of interviews exploring the world of emerging technologies. I'm Paul Thies, your host.

And in this episode of If/When, we will be discussing artificial intelligence and data analytics with Stratos Davlos, the newly named chief technology officer for AI firm Innoplexus and Reuben Maher, Jacobs' [00:01:00] global director and data analytics lead.

In his new role, Stratos will be charged with developing an Innoplexus research lab in Silicon Valley to accelerate its AI development for the firm's drug discovery and development program.

Just prior to this role, Stratos was the vice president of Watson AI and Engineering at IBM, where he led a team of cognitive technologists to create industrial strength AI Watson products, by combining IBM's extensive research and engineering capabilities.

[00:01:30] As the global director and data analytics lead for Jacobs' aerospace technology and nuclear line of business, Reuben is charged with standing up an enterprise wide predictive analytics initiative that incorporates artificial intelligence and machine learning.

His career spans over 20 years of service and counseling governmental and commercial clients on their data strategies. Welcome to both of you. To start us off, I'd like to ask Stratos, what do you see as the top two to three benefits [00:02:00] that artificial intelligence can provide to organizations?

Stratos Davlos: Well, first of all, thank you very much for having us, Paul. In general, AI helps companies to improve their business efficiencies, make them more competitive and show they're in compliance with regulations.

It depends on the line of business. So if you're working with a marketing organization, AI can monitor, for example, publications and articles from competitors and summarize their strategic [00:02:30] intent. You can imagine how beneficial that can be.

If you work in a support organization, AI can improve processes, automate decision-making, automate tasks, behind the scenes tasks. And really the promise is to make their employees more efficient and the customers happier.

If I was to put the spotlight a little bit in the current business I'm involved, in life sciences [00:03:00] drug discovery is an issue, and AI is employed to solve this vital problem. So if you do drug discovery, you go through multiple stages.

It's very expensive, it's lengthy and it's risky. So if you think about the drug candidates and how many of those pass at a clinical trial stage, the success rate is as low as 10% to 15%. And that comes after a three to four year investment [00:03:30] and a significant amount of money that goes into this.

So if you're thinking about Innoplexus, the company that I work for at the moment, they have a purpose built AI platform that makes drug discovery more efficient and more effective. So it uses some machine learning-driven hypothesis. It generates, for example, potential biomarkers based on unbiased AI models that Innoplexus owns.

And these models are built with an [00:04:00] extensive network of biological validated data, and that data comes from both published and unpublished literature. And the outcome is it ensures that the models have this high biomarker prediction accuracy that [inaudible 00:04:21] over the traditional methods.

Paul Thies: And Reuben, I've got the same question for you from where you sit within Jacobs and the aerospace [00:04:30] and technology sector. What do you see and what are your clients seeing as the top two to three benefits that AI can provide to those organizations?

Reuben Maher: Paul, like Stratos said, first of all, thank you for having us. It's a pleasure to be able to speak on really an exciting and rapidly evolving topic. And I think Stratos did a great job of highlighting some of the things that we're also seeing. It frees up employees to focus on higher value tasks, [00:05:00] greater efficiencies, a couple of other things.

We're seeing things like the ability to rapidly and really, at enterprise scale, be able to process extremely large quantities of data, and then be able to provide actionable insights.

A few years back, one of our clients at Jacobs, we developed and deployed within months, not years, a cloud-based analytic capability on top of a high performance compute infrastructure.

And it had the ability to leverage against [00:05:30] AI machine learning models to identify high frequency trade fraud. And then even probably more

importantly than to be able to gather all this data, store it, and then effectively prosecute the nefarious traders.

So another the benefit we're seeing as well though is companies increasingly need to engage with their clients, and artificial intelligence is bringing enhanced capabilities to be able to intelligently interact with customers in a much more real-time nature.

So I'll dive into one [00:06:00] specific area to drive that point home a little bit. In artificial intelligence, there's something called natural language process, or you'll hear it referred to as NLP.

So NLP is basically a substrata of AI that's used to aid computers in understanding how human beings' normal everyday language works. The goal is for the computer to not only gain this understanding, but to drive meaning, and then value from that human language.

People often say, "Well, how's NLP used?" So if you've ever [00:06:30] been overseas and you've had to use Google Translate, that's an example of NLP-based technology. If you've been on the internet and you've had this customer service window suddenly pop up and start offering to assist you and answer any questions you might have, those are increasingly automated chat bots. They're really becoming more and more realistic and surprisingly accurate, leaving you often wondering if you're talking to a human or a machine.

And then there's of course, some of the personal assistant application stuff [00:07:00] like Siri, Okay Google, Cortana, Alexa. Again, NLP. A practical application of this is a few years ago, I led a team that we were asked to support a large industry-leading national intelligence conference and monitoring some of their social media traffic coverage throughout the week.

And so what we did was we leveraged an industry-leading NLP platform, and it had a sentiment analysis tool built in. And we monitored both the real-time feeds as well as deliver daily reports, showing summaries of things [00:07:30] like number of tweets, top social media platforms, generating content, words with sentiment, trending topics.

So what we were able to do is provide them insights into positive versus negative sentiments. Stuff like, "Hey, I loved the keynote speaker this morning," or things like, "Man, lunch was terrible today."

Comments and sentiments that users were providing back that the NLP, again a substrata of AI, was able then to interpret and provide insights. And [00:08:00] basically, these insights were key components of their conference.

They were able to address, mitigate and make the conference experience even better the following year. So when we're looking at the benefit of AI to our

organizations, there's even a ton more than what we've talked about. But these are some of the encapsulations, or I'd say the low-hanging fruit we're seeing that are making a very, very tangible difference to companies nowadays.

Paul Thies: You bring up an interesting point. Obviously, with all the content and all the data that is generated around the world [00:08:30] and through organizations on a daily basis, it begs the question, Stratos. What do you think is the most pressing problem for organizations, in terms of AI development and using that data, the massive amounts of data that can be collected? How do organizations attack some problems that might arise with AI?

Stratos Davlos: Yeah. I was about to say data. Data availability, in fact data that's ready for AI [00:09:00] is one of the biggest stressors for companies. And the reality is you either have a lot of data, more than you can actually use, or you have no data or no permission to use the data you have.

And there's a known study showing that about 90% of the data in the world today were accumulated over the last two years. So if you have data, how can you make sure that you're actually gaining the insights you [00:09:30] need to lead your industry? And that's what companies are going through today.

They're saying, "Okay. How can I leverage the data I have? And then in order to do so, do I need to have access to data scientists or specialists?" And as you know, data scientists are a rare species.

So the other thing that causes a pressing problem is, "How many people do I need to hire, and do they have the right skills? And do they speak the same vocabulary like I do to understand how they're going to process the data [00:10:00] I own?"

Now owning the data is a liability, and you can get hacked. And you know the byproduct of that and the significant marketing hit that you're getting, plus all sorts of legal implications of that.

You may be infringing upon people's privacy, and you're at the mercy of audit and regulations, especially in Europe you see that. And if you end up having a sliver of data that you can use, [00:10:30] then you build models.

And once you have those models, the next level of anxiety is, "Okay. They produce some outcomes. How can I explain those? And when those outcomes are produced and I get sued as a result if I decline a loan because I'm using a machine learning model that was used from the data I have, how can I explain that to a customer or to a judge?"

Machine learning is also open and vulnerable on adversarial attacks [00:11:00] and all these things. So if you just put all of this together, data availability, ability to prepare data for AI, access to data scientists to build some advanced models.

And then once you have those models, do they do the right job? And can I protect myself in case of a lawsuit? Are there things that are the most pressing problems our organizations have in terms of artificial intelligence?

Paul Thies: You [00:11:30] really hit on something there because it's trying to get the workforce skilled up to be able to deal with these kinds of advances. There's a data point I like to quote. It was from a report by the IDC. And they speculated that by the year 2020, the world will generate 44 zettabytes of data.

And a zettabyte, it's one with 21 zeros behind it. That's how much data is getting generated [00:12:00] or will be generated on a global basis. So this next question I've got is for you, Reuben, it's a segue there.

I had seen earlier this summer that Accenture had released a report stating the aerospace and defense industry will be more affected by AI than by any other major emerging technology, with one third of aerospace and defense executives citing AI as having the greatest impact.

And there's an accelerating focus on reskilling their workforce. So [00:12:30] the question for you, Reuben, is what are some of the most common questions your clients are asking about AI?

Reuben Maher: So it's funny. It's such a complex topic, and yet the questions we often get are so simple. I can get questions like, "What is it?" People have seen these futuristic dystopian technology run wild movies like AI, things like iRobot, and they have this perception that may or may not be realistic, and it may not jive with what's currently [00:13:00] out there.

Interestingly, on AI, it's basically essentially a subcomponent of computer science and it focuses on the development of intelligence within the machines. And it's really designed, trying to make these machines more human-like in their decision making.

So at the end of the day, it's how do we get machines to think in much more human nature to do what we talked about in the first question, make things easier, make things faster, more efficient, and let humans really focus on higher value at work?

Now [00:13:30] that general misunderstanding of the basics of what is AI, it leads to other basic questions. Things like questions I've actually had, "Can't you just do some AI and ML and show me what's wrong and how to fix it?"

And it's almost like this. If I just take an algorithm, I throw it at some data, I'll discover the problem and the solution. It doesn't work that way. And I'll pick on the NLP, which I talked a little about earlier. So that's really difficult, let's say you had...

If anybody's had greater than 50% success [00:14:00] with any of the Siris or Alexas, that's really a credit to how far the technology has come. It's difficult because there's so many intricacies and nuances in the NLP case of the human language. For instance, you're dealing with different languages, you're dealing with slang, sarcasm, context, ambiguity, things like misspellings, colloquialisms.

It's funny. An example of the challenge of interpreting just the human language is a humorous incident back in the 1950s, when some words were being translated back and forth between [00:14:30] the English and Russian languages.

And here's the sentence that required translation, "The spirit is willing, but the flesh is weak." Now when that same sentence, which translated to Russian then back to English, here's what they came out with, "The vodka is good, but the meat is rotten."

It just goes to show that translation in general, much more so in a scenario where you're having to teach computers context and colloquialisms and regionality, it's [00:15:00] really, really difficult.

But some of the advances in machine learning, artificial intelligence, compute capacity capability, they're making some of these horizons now the here and now, and not the distant future.

Paul Thies: So Stratos, from where you sit, what do you see are some of the big misconceptions that people entertain about AI? And conversely, what are the realities?

Stratos Davlos: I think Reuben hit it right on the nail. People think AI is magical, "I'm going to [00:15:30] buy this AI computer, connect it to my network and now I will figure out how to optimize my business." And that's far from the truth.

AI really is an ecosystem. It's an ecosystem that includes the computer, the machine, but also the world as it's represented by data, and the business goals that need to be clearly stated, and also the user or the systems that it actually serves.

[00:16:00] So these four elements provide the context for AI to be successful. And it's important to highlight that AI is really a symbiotic relationship between the computer, the machine and the human. And really, the human needs to augment the machine before the machine can augment the human.

Another misconception is that really AI's going to take away our jobs. And really, the mission of AI, from where I stand, is to augment [00:16:30] human capability. I believe that it's going to create new jobs.

So Ginni Rometty, the CEO of IBM, actually calls this a new color of workforce. And it's not just data scientists that are going to get employed and reap the benefits of this new era. There's also any human can participate in this by doing crowdsourcing and labeling and generating data and validating, so to speak, whether AI is doing a good job.

It's not [00:17:00] going to take away our jobs, it's going to generate more jobs. The other thing is people really believe that AI is precise. And it's precise just like the way it is when you code or you use a rule, or the way we do programming today.

And the reality is it's not. It's highly accurate, but it reflects the state of the world. And our world is inaccurate with biases and ambiguity and conflicting [00:17:30] opinions. And it gets really confusing, especially when you're in an NLP part of it, just like the way Reuben explained how translation from English to Russian and back can create a total mess.

I believe that once we demystify these things, and once we explain that we are departing from explicitly programmed systems and we're moving towards systems that are more data-driven and [00:18:00] the behavior is embedded in the data, then we need to invest in the relevant tech to make sure that we meet that expectation.

So to bridge that gap, there are companies like IBM and Google and others that created capabilities that provide explanations. So when you see an outcome, you can understand exactly why that outcome was made by a machine learning model, can automatically detect any bias that might be either in dataset.

[00:18:30] Or because the system is exposed to a new world, the reality world, you can identify any drift between the system design performance and the runtime design performance, so many others. So my advice would be, have the right expectations for AI, use the right tech to ensure that you bridge that gap between the expectations, and that's it.

Paul Thies: I've run into that too. I think that you do see that there's this [00:19:00] mindset that AI is a fix all or people don't really necessarily understand how it works. They have this assumption that it's some sort of magic that will fix things.

But one of the things that's really interesting is there's a bit of fear or anxiety about what AI and automation will bring to the workforce. And in fact, Deloitte just recently released its 2019 Global Human Capital Trends Report, and 64% of the [00:19:30] respondents cited AI and robotics as an important or very important issue in human capital.

But when you look at it, a lot of that impact where there's some disruptive impact is really going to be at the white collar level where there's intelligence

work that, for instance, accounting or financial services or things that can be automated with algorithms.

But something that I found really interesting and Reuben, this hits [00:20:00] home for Jacobs, is that McKinsey Global Institute did a study on occupational shifts attributed to AI. And they found that AI will actually mean employment growth for engineers, architects, infrastructure workers, and others in similar positions.

And in fact, the numbers were really pretty startling. They're expecting an increase in those occupations by 35% in the US by 2030, [00:20:30] and with higher increases in Mexico at 48% and India at 117%.

So AI's not replacing jobs or AI's not replacing work for humans to do, but it's changing the nature of work. Let me ask you, Reuben, what are some of the issues that you are working to solve for and your team at Jacobs using AI and data analytics?

Reuben Maher: Yeah. This is a great question. We've been given a [00:21:00] lot of taskings from clients, a lot of projects from clients that they're wanting to take traditional issues they've had, whether it be manual data processing, whether it be trying to limit the amount of downtime of systems and machinery and factories and plants.

How can we take this data that they already have and make it more intelligent to drive insights, and then make better decisions quickly? A couple quick examples. In North America's hot desert in Arizona, [00:21:30] water's a vital resource. But the problem is it's extremely limited in supply.

So Jacobs, we've partnered with Pima County and we've delivered really a truly advanced water reclamation facility, which helps save taxpayers not only millions of dollars, but also augments their really, really precious and valuable water supplies.

So they had a challenge though. They had some of the aeration blowers that were failing and any time there's failure, that's downtime, and [00:22:00] the whole process of the water reclamation is impacted.

So how can we, using some of the data, predict, avoid failures in these aeration blowers? What our team did was we gathered two years of historical SCADA data... And the SCADA systems are really these computer systems that they gather and they analyze real-time data.

And so we took that data, put it in the cloud, and then fused it with CMMS event data. So these CMMS systems are these computerized maintenance [00:22:30] management systems, and they basically help manage organizations' maintenance operations.

So we also leveraged artificial intelligence and machine learning, specifically a neural network, in the development of a digital twin to help forecast and predict asset failure by running multiple what-if scenarios with the data we had. The benefit was that we were able to end up predicting blower failures 50 days in advance, thus resulting in increased downtime and replacement costs.

Now, a totally different use case [00:23:00] of artificial intelligence being used in real world is for NASA. Jacobs is NASA's largest professional services provider in the world. We do billions of dollars of business with NASA. Everything from managing plants, managing facilities, to things like the International Space Station.

We develop launch control vehicles, 3D printing onboard the International Space Station, as well as experimental things like how do we develop 3D printing materials for potential habitable structures on Mars? [00:23:30] So these are the kind of things we're doing for NASA.

So NASA came to our team and said, "Hey, listen. What we want to do is our mandate is to have boots on the ground or boots on the Moon by 2024. What we'd love to do is we have these millions of images of the moon. We want to be able to take those and leverage artificial intelligence machine learning to help us predict and recommend an optimal lunar landing spot."

So what the team did was started working on the really, really basic building blocks of that. You don't just go and [00:24:00] feed all those to a computer and say, "All right. Recommend it." You've got to train this model from the ground up. It's really the concept of crawl before you walk.

So we developed a cloud-based advanced analytics solution for NASA on Microsoft's Azure artificial intelligence platform, leveraging Azure's Cognitive Services suite, specifically their AI-based computer vision capability.

So what the solution does is it combines facets of AI and machine learning, computer and machine vision, and image processing to automate. And that's one of the key parts [00:24:30] here is the automation and binary categorical classification of city imagery.

So we fed it all these images of cities. So we had to first train the model. So we first ran a binary classification of, look at these thousands of images and come out with a decision on if it's a city or not a city.

After successfully completing that, we then went on to categorical classification. "All right, now which city is it? Is it London? Is it Paris? Is it Rome? Is it Sydney? Is it New York?" And then after that, it was basically taking these images of cities and [00:25:00] geo-placing them on a map, almost like a puzzle piece.

What we found was by using this artificial intelligence capability, we were able to fast track our team's capability via automated tools instead of manual processes, to leverage artificial intelligence and to search and analyze these massive amounts of images faster and with higher fidelity.

It really resulted in higher accuracy rates, much more quickly trained models and greater object recognition. And this was a huge win for our NASA client. They've actually taken this out of proof [00:25:30] of concept and have funded it and moved it forward.

Paul Thies: That is amazing. Can you imagine the Apollo 11 team having to pore over a million images of the moon to be able to locate where they would put the Eagle down, for instance? It's really amazing.

Going back to the idea of we've got so much data and so much content out there now, it's almost impossible for humans, even a collection of humans, to be able to really efficiently [00:26:00] pore through that. That's where that idea of the magic of AI comes in as being able to do that.

Stratos, let me ask you. Reuben was sharing some of those amazing things that AI is able to do for Jacobs. What do you see where you were sitting with IBM, and then also where you're going with Innoplexus? What do you see as the most significant element of AI development in the next few years?

Stratos Davlos: Well, I'll tell you what surprises me with AI. [00:26:30] To date, we made a lot of progress in the field of AI. There's no doubt about it. In particular, if you look at the deep neural networks, the recurrent networks and some of the sophistication behind [inaudible 00:26:45] and others, there's a lot of innovation that goes into those. And there is a really passionate and enthusiastic talent pool that innovate every day.

The thing that surprises me that is the opposing forces [00:27:00] on AI. AI is vulnerable to attacks. So robustness is what we call for AI is a big issue. To my surprise, there's a concerted effort out there that's trying to manipulate outcomes at one time, in various forms.

So you can poison the training data. So if you use data out in the wild or the data a company collects has vulnerabilities [00:27:30] at the collection point, then you can manipulate the outcomes at runtime. So an adversary can do that.

They can also contaminate theaters and inject noise, either audio noise or visual noise, in a scene. So if you have a self-driving vehicle, the vehicle stops seeing pedestrians crossing the road by just inserting sticky notes on the scene that the vehicle sees.

What I'm saying right now to [00:28:00] you, if you had an AI system that transcribes it, somebody can play white noise that can confuse the transcription system to have a completely different outcome to what I'm saying.

So these are the things that caught me by surprise, and I got a lot more exposure to them while I was at IBM. And there is a big group at IBM that's actually looking around corners and trying to mitigate those [00:28:30] concerns.

Paul Thies: Yeah. And one of the things that I'm interested in is organizations and how they can share data. Obviously with cybersecurity and data privacy concerns, it becomes a very tangled interchange. But how organizations would be able to share data to empower their AI to accelerate that development but at the same time, there can be risks. It's fraught with risks, geopolitical risks, depending on who the players are.

Stratos Davlos: [00:29:00] I can add some color to that, if you like.

Paul Thies: Yeah. Yeah.

Stratos Davlos: Where we're heading towards is AI models that are built without data, by creating local models using the data you have access to, and then sharing the corresponding gradients in either a differential private gradient or [inaudible 00:29:25] encrypted gradients with other models, and [00:29:30] then have a mechanism of merging those gradients. It's called federated learning.

And I think federated learning is receiving a lot more attention in our days. TensorFlow released a capability like this. IBM has numerous papers and assets in that area. At Innoplexus, we are using federated learning to ensure that sensitive clinical trial data that happen in different regions but different countries are not leaving either corporate [00:30:00] walls or those regions, for regulatory purposes.

So how do you build the insights when you're doing drug discovery, drug purposing using the global data sets? So federated learning is a technique that we're employing at Innoplexus to make that happen.

Paul Thies: And Reuben, I can imagine that you're all at Jacobs dealing with aerospace and defense clients, and there's a lot of sensitive information that has to be handled. I know that also part of your role is standing up [00:30:30] innovation work streams within Jacobs. How do you see organizations using AI and data analytics in support of their innovation activities?

Reuben Maher: So it's interesting that when we look at bringing innovation into companies, people will often look at this as a function of technology or a function of people, and they'll try to solve it from that perspective.

Now, I'll caveat my answer with we are not in a technology business. We are not in an engineering business. [00:31:00] We are not in a data business. At the end of the day, we are in a people business. This is all about people and how we interact with them. How we engage with them drives pretty much the solutions of everything we do.

That being said, the challenge of innovation is not simply a technology or a people problem, although those are factors. I think that the overarching key to achieving innovation across an organization really is integrated into when you solve the cultural problem and you solve the cultural challenge.

So things like artificial [00:31:30] intelligence, machine learning, cloud computing, IoT, cybersecurity, they often run counter to how non-digital companies think and operate. But how does a company effectively change its culture?

First, there's got to be recognition that this is difficult. It's not easy. And it's going to take intentionality and top down leadership. As a result, it's also going to take significant investments of time and resources.

And we have to avoid the expectation that this is going to happen in weeks or even months. [00:32:00] It's going to take longer. I think the second piece is innovation. It's really going to happen when organizations create an environment where employees are encouraged to fail and learn quickly, and then be able to pivot, and then in an agile manner, iterate.

Now, it's got to be counterbalanced with... We've got to have an intolerance for incompetence. If individual or teams, if they're not learning from their mistakes and getting better actively, then corrective actions and course corrections have to be taken, and quickly.

There's also got to be, I think, some constant [00:32:30] dialogues that encourage not only open communication but on the other hand, well, very candid feedback. And I remember, and I've got something here from a Harvard Business Review, they encapsulated this really, really well in their January, February, 2019 article.

And they stated, and I'll quote it, "When it comes to innovation, the candid organization will outperform the nice one every time. The latter, meaning the nice organization, they confuse politeness and niceness with respect. There's nothing inconsistent [00:33:00] about being frank and respectful. In fact, I would argue that providing and accepting frank criticism is one of the hallmarks of respect. Accepting a devastating critique of your idea is possible, only if you respect the opinion of the person providing that feedback."

There's got to be this back and forth communication that keeps these lines of dialog open. Accountability has got to be central to this. So we've got to be able

to understand and really clearly articulate the business problems we're trying to solve.

But that being said, what [00:33:30] you can't measure, you can't manage. So there's got to be metrics that are being gathered to determine how well the organization is performing. And there's got to be some type of regular review process that ensures progress against whatever the stated goals are, these innovation goals, that they're being met or that return on investment metrics are being realized.

And I think it's important to restate, we don't want to be in the business of developing solutions in search of problems. We need a strong business case. We need that willingness to place big bets.

We need [00:34:00] an agile development type of mindset, and then some type of gatekeeper accountability model to know when to exit and when to go full speed ahead towards building and deploying some of these innovative technologies like artificial intelligence and machine learning.

Paul Thies: Stratos, let me ask you this. Let me shift gears just a little bit here, but we were talking earlier and you had recently left your executive position at IBM and the AI team and you joined Innoplexus. And what informed your decision to join the Innoplexus team?

Stratos Davlos: [00:34:30] The first thing to say is IBM is an awesome company. I had the honor and privilege to work with really some of the best AI scientists in the world. So the question is why Innoplexus? It's personal.

Over the last two years, I had cancer hitting my family, my immediate family. Four times. One resulted into a fatal outcome. And when that happens, you [00:35:00] feel hopeless. There's no money in the world that can buy the cure.

And when you actually look at the drug development business, it's painfully slow, expensive, ineffective, and something needs to be done to accelerate this. And Innoplexus is solving that problem.

So to me, besides the coolness of the AI technology they have, they're actually serving a purpose. [00:35:30] And that's why I thought it's time for me to shift my energy into something that's very, very meaningful and makes a dent in the universe and helps humanity. So that's the reason I chose Innoplexus.

Paul Thies: Thank you for sharing that. I know you're pretty excited about this next chapter in your career, and it's a new opportunity. And obviously, when you can marry that to something that's personally meaningful and it really hits home, [00:36:00] that's a real blessing. So thank you for sharing that.

Our last question here for today is for both of you. And Reuben, we'll start with you. And then Stratos, I'll ask you the same question. But Reuben, what would you say is the most surprising thing you've learned working in AI and NLP and data analytics and in this whole field? What's the most surprising thing you've learned?

Reuben Maher: Earlier, we talked a little bit about how complex the applications of advanced analytics, some of their capabilities like artificial [00:36:30] intelligence and machine learning are. And the underlying algorithms and computational procedures, they require incredibly skilled analytics [inaudible 00:36:40] to really fully develop and enhance and tweak them in order to provide much greater value for very specific challenges.

However, we're also seeing that these really advanced technologies packaged and delivered to consumers are providing these really, really accessible artificial intelligence [00:37:00] solutions. You're increasingly seeing easy to use interfaces with surprisingly effective results. I'll give you a couple examples.

So just this week, Tableau, who's one of the industry leading business intelligence and data analytic providers, they announced a new feature called Explain Data. It uses AI to automatically unearth insights more rapidly.

So what it does is it takes the data that they're already working with and it saves these users valuable time they'd typically spend manually sorting through [00:37:30] mountains and reams of data. And instead, it's continually performing analysis for them in an automated fashion.

And it's still very, very new. But what they're claiming is that Explain Data also helps remove bias in the data while providing a more rigorous outlook of the data. And the bias is a really significant problem in a lot of the industries nowadays [inaudible 00:37:51] with data.

And Stratos talked a little bit about it earlier. And this is something where as we're looking to... You talked about federating some of the learning aspects and federating some of the [00:38:00] data processing.

We've got to make sure that the underlying data is unbiased and it's producing results that actually are relevant to broad swaths of the population, not just to specific niches. So, additionally, I think Microsoft Office, they've...

Surprisingly, these are tools that pretty much everyone uses on an almost daily basis, but they're increasingly incorporating artificial intelligence into these tools. For instance, at the Microsoft Inspire conference earlier this year, they introduced an AI-powered feature [00:38:30] in PowerPoint called Presenter Coach.

And when they showed it, it was really, really cool. So it uses machine learning to coach you, and then it delivers a report with tips and suggestions after you've been practicing. And it helps you deliver a more powerful presentation.

If you're reading too fast, it's going to recommend you slow down your pace. It also incorporates some kind of contextual understanding as well. So if you say, "Hey, you're the best man for the job," it'll inform you that this may actually be culturally insensitive [00:39:00] and instead recommend an alternate phrasing such as, "Hey, you're the best person for the job."

Likewise, in Excel, they've added a capability. It's right there on the little home tab, and it's a little button called Ideas. And you click on it and it takes your spreadsheet and highlights patterns in your data and brings out new visualizations and insights.

And it can identify trend patterns within time series data, things like your costs have increased or decreased over time, or it can identify where a majority of total value is due to simple [00:39:30] factors. Things like teenagers account for the majority of your cell phone bill.

And we're seeing these really difficult and complex things under the covers be presented in a very simple and easy-to-consume manner for users. We want to be in the business of building solutions, and not just solutions in search of a problem.

That's where I think it's exciting to look at some of the things that we're seeing in artificial intelligence and machine learning because these are really functions of problem solving, insight generation.

It's important that we continue. We've got to continue to keep our clients focused [00:40:00] on these bigger pictures of what's your business or mission problem you're trying to solve, and then what's the best solution to solve it? And then turn those into really practical solutions that users can quickly and easily incorporate and use to provide value.

Stratos Davlos: I promised to add a little bit more color to my answer as to what surprised me. I mentioned two things. One was the energy and effort going into AI. And the [00:40:30] other one was how vulnerable AI is and the opposing forces with adversarial attacks.

Now, when it comes down to the energy we're putting into AI, we actually have a lot of things evolving towards more autonomous systems. So most of the models that we know of are what we call narrow AI. They're like a single task, single domain, superhuman accuracy and speed on certain things.

There are a lot [00:41:00] of models that today have moved from this narrow AI concept to more of a broad AI capability where they're multi-task, multi-domain distributed and explainable and multi-modal, and those models exist today.

So the ability to really analyze an image or look at a video and identify the most important scenes out of, let's say, a soccer game, these are models [00:41:30] that today... And to my surprise, are fairly robust. There are also models that are more towards what we call general AI.

These are these cross-domain systems that can learn and reason with broader autonomy. And some of those that I got exposed to are fascinating. So to me, it's amazing how quickly AI is evolving, and it's maybe one of the most impressive [00:42:00] and surprising elements that I've seen.

Paul Thies: Stratos and Reuben, I want to thank you both for joining me today and sharing your insights. We've just started to scratch the surface of complexity and the energy that exists out there in AI development, and what it can mean, and what it means for the human workforce, and what it means for organizations. So I want to thank you both for joining me today and...

Reuben Maher: Thank you, Paul.

Paul Thies: Thank you. And I look forward to continuing [00:42:30] this discussion and seeing more on what's going on in the world of artificial intelligence.

Stratos Davlos: Our pleasure.

Reuben Maher: Thanks, Paul. (silence)