Improving Business Operations With Machine Learning and Artificial Intelligence

You shouldn't need superpowers to complete a mountain of work or find answers to complicated and confusing questions. The advent of artificial intelligence (AI) and machine learning (ML) has generated the ability to put time-consuming, labor-intensive, and detail-oriented processes on autopilot — and to unravel seemingly contradictory questions in the blink of an eye.

This webinar explores the practical applications of AI and ML in practical business settings.

Find additional resources on our event page: https://www.claconnect.com/events/2022/improving-business-operations-with-machine-learning

Here is a transcription of this session:

Troy Hollings:

Okay, thank you everybody. My name is Troy Hollings and I will be our host and excited to be here. Thanks everybody for tuning in, looks like we've got about 253 people right now. So welcome everybody. How it is my distinct honor to be bringing a guest speaker here today, Spencer Lourens. He's our managing principle of data science, machine learning AI and fun fact, that headshot, it's a new headshot. He looks great. He grew up on a farm in Central Iowa, PhD in Biostatistics and graduated in 2015. Briefly was a Professor and then basically parlayed his statistical knowledge into all things machine learning, joined CLA in 2019 and has experience with non-linear machine learning algorithms, random forests, boosting, SVM, I don't even know what that means. We'll talk about it today. Neural networks and more. Rivaling only Thomas Edison in raw intellectual horsepower, he's also a super nice guy. Spence, welcome to the show. I think we can call it a show.

Spencer Lourens:

Thank you so much, Troy. I don't know if I rival Thomas Edison at all, but I appreciate that.

Troy Hollings:

Anything you want to add on your distinguished bio before we jump in?

Spencer Lourens:

Oh no, you did a great job. Thank you so much. It's awesome.

Troy Hollings:

Appreciate it. Okay. Let me set a little bit of context here, what are we going to talk about? So I think there's a lot of confusion in general, about machine learning. What is it? AI, is it on one end, a panacea that's going to solve all problems? And on the other, Ray Kurzweil said that it's the end of the world and the world's going to end in the singularity with the AI being one of the largest risks of human race. What is it? Spence is going to... Is here to teach us the way. So what we're going to talk about a little bit more tactically here, little bit of definitions, brief history of AI, machine learning, all of that, some well known applications. If we have time, go into some specific use cases and then just even thinking through how could you as a private company, as a government agency, as a school, how could you be using AI or machine learning?

Troy Hollings:

So let's jump in. What is it? So Spencer, as I'm sure you know, Socrates said, "The beginning of wisdom is the definition of terms." So let's define some terms. I see deep learning, machine learning, AI, you mind giving us just the overview on that?

Spencer Lourens:

For sure. Yeah. I think what I'll do is I'll start from the right and go to the left. So I'll start with artificial intelligence. Again, like Troy said, these three terms get conflated a lot or just used synonymously, which really is fine for the most part. But today we're going to try to give you a little bit more of the signal and the actual different definitions. They're listed up there on the slide, but essentially artificial intelligence is taking a machine or a deep learning model and actually deploying it in a framework or in some type of application so that it can make decisions and try to mimic what you might see humans doing as making decisions. So for instance, when we're little kids, we learn to crawl, we learn to walk, we learn to pick things up and carry them. We're at the stage today where robots, if you've ever heard of Boston Dynamics, can learn to walk over a series about 10 minutes through just trying to move, trying to do different things, make different actions. That's really what artificial intelligence is.

Spencer Lourens:

It's equipping a computer program or some type of agent with a machine learning algorithm, so that they can mimic human intelligence. So that's AI, that's artificial intelligence, machine learning is second in the middle and it's really the algorithms that make up what we do when we're engaging in artificial intelligence. So it's taking historical data. It's taking signals from the environment in which an agent or an algorithm or a computer is sitting and understanding what happens when I make certain choices or I make certain actions. It can also be, which is something we're going to show later in this call, is something like, which ones of your customers are showing signals that they're likely to leave your client base? Which ones of your employees are showing signals that they're not as engaged as you want them to be? And they might be getting ready to leave your workforce.

Spencer Lourens:

So that's really ML, it's kind of a subclass of AI, but it's going to be deployed in an AI system when you're actually getting to that stage. And finally, deep learning, at least in my field, this has been a huge buzzword. What makes learning deep? What is deep learning? What does that even mean? It's basically a subset of machine learning where we take something called a neural network, which is essentially just a really complex matrix multiplication if you want to think about it that way, I know that's a little cerebral, that's a little hard to understand if you don't come from a math background. But think of it as a newer algorithm, at least in terms of how popularized it is. And it's very deep in terms of it has a bunch of these things called layers and it can accomplish really, really complex tasks.

Spencer Lourens:

So things today like if you are taking a picture of your ID or a credit card for your bank, if you've seen how they can scan information out of that, out of just taking your credit card and listing it up, most likely that's using what's called a deep learning algorithm. All of these algorithms today about taking images or videos and mining them and automatically taking information out of them, are basically going to be using deep learning under the hood. So that's what deep learning is. So those are the three... I'll show a couple plots or slides on deep learning later on that hopefully will kind of edge out some of those little... Those rough edges in my explanation. But yeah, those are the three. Those are kind of the things we're going to talk about today. I think next we're going to jump into the history of AI, which has been a little sparse in like a long time ago, but it has certainly filled up since 2015 until now. Anything you want to add, Troy?

Troy Hollings:

Yeah. Well, I think that there's a little bit of a misconception that AI is this new thing, that it just started in 2015 or 2020, or it's this new thing, but it's been around for a while. So I'd love it if you'd tell the story of AI, feel free to include any interesting facts, interesting fact I learned is that there's actually a robot in 2017 who got citizenship to Saudi Arabia, didn't know if you knew that.

Spencer Lourens:

Wow. No, I actually didn't. That's awesome. Maybe that will happen more often in the future. That's a little scary for some people probably. Cool. So as Troy had mentioned, AI has been around at least as a concept for a long time. On the timeline from left to right, you can see that John McCarthy originally coined that term for just a seminar at a conference in 1956, this idea of artificial intelligence and what that would be. And of course, we know that Turing played a big role in this. There were a lot of computer scientists that were involved in this early on, and they did some pretty fantastic things. But in 1960, perceptrons were around before 1960, but 1960 was really the first time there was evidence of what's called a perceptron, which is really a foundational building block of neural networks.

Spencer Lourens:

It's a first time that it learned by trial and error, i.e it randomly tried to predict something, it realized it was really bad at it and then it updated its predictions, realized it was a little better at it, it updated its predictions more, got better at it. It kept just learning from feedback that was pushed back and forth through the system. And it learned how to do a task better, which that could just be predicting somebody's height based on their father and mother's height, their age and other factors, that maybe their weight, et cetera. In 1976, we saw the first automation of a robot actually being able to interact with assembly lines to increase production rates. 1997 is really exciting because that's the first time that we saw a robot be able to beat a human as something that up to that time, it basically seemed impossible.

Spencer Lourens:

Chess is a complex game. I mean, I'm not good at it by any means. I probably lost a lot more chess games than I've ever won. But the idea that a computer could beat somebody that was a master at that, was basically... It was a dream, right? And in 1997, that actually happened. We saw just how flawed AI was after 1997 though. And we kind of fell into what's called, I think it was the third winter of AI, where AI has over promised a lot in the past, especially in the 1960s, 1980s, where people were like, Oh, it's going to do these amazing things. It's going to be able to drive automatically. It's going to be able to do things that we're actually seeing it really do today and it didn't deliver. So that's what kind of happened after 1997, is we went into another winter, but then finally 2011, some of these things in deep learning started to happen. We started getting more computing power. We started to get access to more powerful machines. So Apple introduced Siri, where you could actually say, call Troy Hollings, set up a calendar meeting with Troy Hollings to talk about Digital at CLA Saturday at 8:00 AM. He's going to decline that meeting.

Troy Hollings:

Yeah. I'm going to miss that, yeah.

Spencer Lourens:

But you can actually do that, right? So that's an idea of voice to text. So the machine is able to actually listen to what we're saying, take those sound frequencies and that data that's produced and translate it into text. And you can probably see today, if you do text... What is it? Speech to text, a lot of people like to do, not always perfect, but it's at least doing pretty well in terms of getting what you're trying to say, translated. Maybe you just have to be a little bit more careful in terms of the speed at which you talk or the way you enunciate.

Spencer Lourens:

In 2015, we had... This is probably where some of the fear starts to get here honestly, Troy. Alexa and Virtual Assistance and I think a lot of people have stories about, Hey, I'm just hanging out with my family, Alexa, stop listening to me. No. No, no, no, don't do that or I have stories on my own where I'm talking with my wife about something I want to buy or something I want to do, I get on YouTube to look up more information five minutes later and it's one of the first options. Honestly, this can't be coincidence, you know what I mean? So there is that listening. So I think one of the things we want to talk about is this is true, how does CLA use AI later today, right? When we get into some of these examples and how we deploy AI solutions, because there is that fear around these systems are listening to me, they're collecting data I don't want them to collect, is that what CLA is saying they're doing, right? So I think we should definitely-

Troy Hollings:

Yeah. That's a good point.

Spencer Lourens:

We definitely want to talk about that, right?

Troy Hollings:

But before that, because I know that this is a safe space, we're in the company of over 300 friends, we're going to take a little risk. We're going to do an activity here. And so the goal of buzzword bonanza is to stumble the PhD. And so audience, I'm going to need your help. Our goal is to beat Spence. And so what we're trying to do is which of these terms is the hardest? Which of these do you think that Spence just doesn't know the answer to? And there should be a survey popping up. If you'd vote, we're going to make him define it off the top of his head, which one do you think that he can't define? I'll give it 30 seconds or something. Oh, oh, okay. Coming in strong. Deep learning so far seems to be the one that they want more definition on.

Spencer Lourens:

Okay.

Troy Hollings:

Here we go. It's still going.

Spencer Lourens:

Okay.

Troy Hollings:

I think we'll give it 10 more seconds and I'm going to make the executive decision. Deep learning's close though. Deep learning.

Spencer Lourens:

Yeah. Data analytics' popping up there.

Troy Hollings:

Okay.

Spencer Lourens:

Look at this participation. We're getting great participation.

Troy Hollings:

Yeah. Definitely.

Spencer Lourens:

Thank you for participating. That's awesome.

Troy Hollings:

Okay. So Spence, I think the people have spoken.

Spencer Lourens:

Okay.

Troy Hollings:

Off the top of your head, need a perfect definition of deep learning or we all win.

Spencer Lourens:

Yeah. Okay. In terms of deep learning, this is a field in and of itself. It's used for extracting information from text, it's used for extracting information from images, from videos, et cetera. Deep learning essentially means I have an input... I've got to somehow take my data and I've got to transform it or transpose it into a numerical representation so I can feed it through, sometimes honesty, hundreds of layers with billions of parameters. I'm going to do this by example, because deep learning is so big, I think that's the best way to do it. And I know you don't want me to talk for 30 minutes, Troy. So you're going to cut me off if I don't do this quickly, but something like GPT-3, if anybody's heard of it, if you have not, I suggest that you look it up on the internet. That is a deep learning model created... It was Google or DeepMind, which maybe deep mind was acquired by Google at this stage.

Spencer Lourens:

But it is able to do many things, i.e you can ask it a question, it will respond. You can have a conversation with it and it's shown itself able to have reasonable conversations on many different topics. You can give it text and have it generate more text. So it can actually automate writing news articles. And humans have been shown not to really be able to determine whether it was written by this deep learning model or it's written by a human. We've not done very well at discriminating between the two. So that's what I'd say deep learning is, these humongous models that are so high capacity that they're really even hard to conceptualize for us exactly what they're doing. They're based on principles. And I'm going to show an example of that in a couple slides. That's what I would call deep learning. So I hope you all are happy with my definition.

Troy Hollings:

Yeah. I think the odds might be stacked against us a little bit. That seemed pretty good, but I don't know if that's more of a comment that our average article is really bad or if that the deep learning's really good, but regardless, I think that was a good answer.

Spencer Lourens:

Good point.

Troy Hollings:

One to zero, you're up. So I actually also have affiliation with the farm. You grew up on a farm. I married into the farm life. I married a farm girl. We've got 25 acres, couple horses, some goats. All my neighbors are farmers, so they actually just think I'm unemployed because I work remote and they're like, "Oh, work remote. Is that what we call it these days?" But I have to say, I'm going to speak in the vernacular of Quincy, Indiana, where I live, what in the good gracious is this?

Spencer Lourens:

Right. And you want a definition for that? Is that what you're saying?

Troy Hollings:

I mean, I feel like I probably need some explanation. Well, would you just shine some light on what this slide is? Because I'm not too clear on that.

Spencer Lourens:

Absolutely. Got it. Okay. This is a representation of a very simple, deep learning model. So I want to give you an example here where we're trying to predict somebody's height, all right? You see X1, X2, XM there, what in the heck is that? X1 could be the mother's height, X2 could be the father's height and maybe XM is just the age. So I've got those three things and what the deep learning model does, I'm going from left to right here guys, just so you know. Okay. So there's a one then the person's mother's height, the person's father's height and the person's age, that's four rows...

Troy Hollings:

And this data would be like genome data, like 23andMe could do something like this or?

Spencer Lourens:

It definitely could be, or it could just be that I did a study of people who are 17 years old and I collected that data from them. I said, "How tall is your mother? How tall is your father?" I tried my best to confirm that. I took their age down. And then what I tried to do is predict their height. That's what I might do. So what a deep learning model does, believe it or not, is you see those weights, W0, W1. W2, W3, in the red? You see that? The next one over?

Troy Hollings:

Yep.

Spencer Lourens:

It just says, okay, I need to wait the last factor, which is a person age by W3, so multiply their age by W3. I need to multiply their father's height by W2, their mother's height by W1 and I'm going to add W0, I'm going to sum that stuff up and I'm going to throw it into some function and I'm going to see how well do I predict on average people's heights from that data. So from their mother's height, their father's height and their age, how well do I do? And I talked about it, iteratively, it learns. So what it does is say how close on average did I get? It uses something called back propagation. That's where that model update comes or that weight update comes in and it says, let's do a little better, it updates the weights. Does it again. Figures out how well it did, does it again. It keeps doing this until it's gotten close enough on average, or it can't get any better on a successive step then it says, okay, I think I now have the best model that I can for predicting somebody's height from their age, their father's height and their mother's height.

Troy Hollings:

I assume the size of the data set probably matters too. If you have a hundred people versus all the data of 23andMe. Is there diminishing returns? Maybe talk about that for a second.

Spencer Lourens:

For this kind of model, because I'm only putting three things in, there's probably at some point where I can't do anymore. I'm only putting three things in, so why do I need 10 million observations? Something like maybe a thousand is plenty, but if you go onto that next slide...

Troy Hollings:

I was going to need the translation on that too. It's like Google translate Spence, translate.

Spencer Lourens:

Good luck, right? So, this was just to show what's happened since 2018 or late 2017. There was a paper published, guys, called attention is all you need. And it's everywhere today. Microsoft's models, Google's models, GPT-3, the one I mentioned earlier. BERT. Probably doesn't mean much to a lot of people that aren't in the field on this call, but there's a question about predicting and automating text that we got and BERT is very popular for that. Not as much as it used to be, but this right here, which I'm not going to go through the details, because it wouldn't really accomplish our goal for this call. I think people would start logging off.

Troy Hollings:

That's a very nice way of saying we're not smart enough except-

Spencer Lourens:

No, it's too boring. And actually I'm not prepared to.

Troy Hollings:

Okay. Oh, stumped the PhD.

Spencer Lourens:

You did stump me. If you asked me about that, you would have me stump. But they found that if I'm trying to predict in specific text, the words that happened before and after and a word in particular are important. And we all know this intuitively. "I really love this job but" versus "I really love this job." Really love and this predicts this job. And so, oh, the job's great but if I don't consider the but that happened after job, I'll think it's more positive than it is. That's what this is all based on, is to say I need to pay attention to specific parts of my input, differently. And that's what attention mechanism means. What should I be paying attention to be able to predict better? So very transformative. It just happened. But a lot of the things that I'm going to show, well, the example I'm going to show on computer vision earlier or later in this slide deck and also some of the well known applications are using transformers today.

Troy Hollings:

Great. I think that was relatively clear. My feeble mind's trying to keep up. So if we move to maybe a little bit more broadly, some well known applications and actually as I was preparing for this on this slide, I got a call and it said, "Spam risk." And I thought, huh, that's interesting because that is machine learning, that is AI trying to say, okay, well, based on some calculation, it's a high likelihood that Troy doesn't want to talk to this person. So didn't call him back. But feel free to read them or if you have any comments, anything that you want to just call out with this well known applications of AI.

Spencer Lourens:

I think what's cool about this slide is that probably most people here have been impacted by this. Text editing happens every time that I'm doing a text on my phone. It happens when I'm writing an email. Grammarly is very helpful and uses AI methods, if you're using Grammarly. We see social media... Anybody who's on social media, it's trying to predict or present you things that are similar to what you've already liked. Kind of like a recommendation system. I want to call out a cool experience I had at Microsoft training center when I was last there, I think in 2019 in Minneapolis. They had just a monitor that was sitting in their hallway. And me and my colleagues who were there could go up to it and smile and act angry and make all these different faces and it would pick up with a confidence level, what kind of face we are showing. It would present it in real time on the monitor.

Spencer Lourens:

And what it did is it picked up where are the corners of my mouth? Where's the center of my mouth? Where's my nose? Corners of my eyes. They get really good through these deep learning models. Again, this is going to be deep learning that's applied. To be able to take any image and however many faces are on that image, they can detect what they call facial landmarks and then they can tie those landmarks position to what type of facial position or expression the person is using. Yeah, exactly. So that's what is happening in face recognition today, which is really, really interesting. The rest, I mean, I think we've already talked about chatbots, et cetera.

Spencer Lourens:

So these are basically things that I think all of us have been touched by, not necessarily what we are going to be talking about, but I think it kind of closes a gap in terms of that confusing slide on transformers. This is how complex this stuff is, but it's impacting us. It's here every day. So that's kind of how I like to relate to it, rather than that transformer that is difficult to understand is like you have to read this paper 10 times to even have somewhat of an understanding. And that's truly how many times I'd have to read that paper to understand it. I'd have to read and read and read.

Troy Hollings:

We need the folks like you, who can go in and build the jet engine, but I just want to fly on the plane.

Spencer Lourens:

Absolutely.

Troy Hollings:

And so If we're going to businesses now, I think it's good to have the context, the explanation, but it's like the plane, I don't need to know how the engine works. I just need to know how can I best use the plane. So it seems like we got a couple comments higher level on how AI is being used in businesses. It seems like there's a couple different categories. You want to talk on these at all, Spence?

Spencer Lourens:

Absolutely. Are you talking about a specific question that we have, or are you talking about the slide?

Troy Hollings:

I'm seeing just category automation, data analytics, natural language processing.

Spencer Lourens:

Got it, buddy.

Troy Hollings:

Seems like those are some large buckets.

Spencer Lourens:

Exactly. So automation enhanced with AI is one of the buckets that we generally attack or see businesses using AI today. So it may be something like I have invoices, expense reimbursements, other type of documents coming in and I want to extract data in some type of structured fashion out of them. And I don't want somebody to have to do that manually. I want a computer to help me do that. So that's really part of this. The dangerous tasks would probably come down to things that need to be done in a manufacturing capacity, where there's risk of injury in an area of the floor where manufacturing is taking place, but there's still use for some type of discrimination or something to happen. So we can remove humans from having to do that so that a computer can do it. Data analytics-

Troy Hollings:

We'll talk about this, but I think one of the things that I'd encourage everybody to keep in mind is that this AI is trying to make people go from being doers to reviewers. And so we have a slide on that, we'll chat through that. But when Spence said that, a light bulb went off in my head like, okay, that makes sense. We're still wanting that human oversight. We're wanting to be able to look at it, to approve, we're not saying build the robots to make all the decisions. It's more like, what are those tasks that are either dangerous, that are repetitive, that there's really no value in a human doing if we can automate? So cut you off, bud-

Spencer Lourens:

Right. And a lot of people wouldn't do it, right? It could lead to a lack of engagement, a lack of satisfaction in the workplace for the individual that has to do that day in and day out, which can lead to turnover cost. So yeah, just backing up what you were saying, Troy. I agree.

Troy Hollings:

And then maybe if we can skip the data analytics, natural language processing, I think we've talked about that. I think that pretty interesting though, if we just... Let's talk about some example uses for AI in business, we've got this slide, but I am committed to stumping the PhD. So I got to ask, over the last year or so, what's been the most interesting way that you've seen a business or a paper you've read or a government entity use AI?

Spencer Lourens:

Yeah. The thing I'm most interested in today is extracting information from unstructured data sources. So we have a slide about that, where we can take an invoice that comes in and apply what's called transfer learning to be able to pull... Here it is. Information that would be meaningful for a process later down the line. So that's probably the most exciting thing that I've seen and it's something that we've been experimenting with and using. And it's one of our applications of actual use cases that we talk to clients about today.

Troy Hollings:

Let's just jump into that if you're okay with it. So I think this is super interesting and I think everybody kind of intuitively understands the obvious, low hanging fruit ways that automation can happen. There's a lot of work to be done there, but I think, let me summarize and you correct me, but one of the challenges in general is PDFs. So PDFs come in and right now, in everywhere basically, a person has to look at that. It's not an Excel field where you can go to the number, you can pull into that number and then use that data. You're having to look at it and we see that actually I spent a lot of my time in Higher Ed and the endowment accounting process.

Troy Hollings:

We see a lot of folks that have got these endowments with a hundred funds, then they've got fund driver, they've got their general ledger, but there's PDFs coming of statements and humans having to go in and click into it. And so if I'm understanding right, this is basically, you build the model that can go in and use that intelligence in the PDF to say, "Oh, I recognize, this looks like a receipt." You want to clear that up or add anything to that?

Spencer Lourens:

No, that's absolutely right. But the bounds of that application are really broad, i.e it may be as simple as I need to classify documents into other categories, which ones are expense reimbursements? Which ones are payables? Which ones are other invoices? Which ones are statements of work? All kinds of things that might be filtering into an inbox, where we can actually hook an automation up to that to extract attachments and then classify them and send them on to other parts of your workflow downstream. So it could be even less targeted in terms of extracting the information like the line by line, which is what this is doing. It's actually maybe a little hard for some of you to see, but it's actually extracting who the bill to is. It's making this image queryable, is what it's doing.

Spencer Lourens:

It's saying the bill to is that accounting firm, it's saying what the invoice number is, the invoice date, all of those things, it's actually providing that from an image. This is not a structured document, it's just an image. And it's able to do that off the cuff. We're also able to set it up so that if you weren't happy with what it's doing, you can make modifications. So again, the human's the reviewer, instead of the doer, that's the way I really see AI.

Troy Hollings:

Can I jump in and ask a clarification on that, Spence? So I think there's always going to be things that, well, I don't want to say always, I can't predict the future, but for the next foreseeable future, there are going to be things just can't be automated. But what you're saying is maybe run your process through this and automate as much as possible. But even if like in that endowment in accounting example, there's a hundred different fund companies and right now that person's logging into each fund, exporting to PDF, looking at it, typing into Excel, maybe there's four macro level categories of those types of funds. Maybe it's bond funds, maybe it's... I'm going to quickly reveal myself here. But what you're saying is build that automation in there as the workflow to, maybe there's still that human doing some of it, but it's way easier to start with four categories and the data mostly inputted, as opposed to just starting at square one. Is that what you're saying?

Spencer Lourens:

Absolutely. Yeah. We can utilize AI to just help increase efficiency even more than creating an application that does the automation part. Absolutely.

Troy Hollings:

Okay.

Spencer Lourens:

So I do [inaudible 00:29:38].

Troy Hollings:

Go ahead.

Spencer Lourens:

Oh, go ahead.

Troy Hollings:

No, no, no, no, you go. I was going to take a different direction.

Spencer Lourens:

Well, I was wondering if you could go back to slide 16, because I think that's another piece that's going to be relevant to people who are here that want to understand how we can apply AI and ML for operational. So there's kind of two buckets here that we put these things into when it comes to use cases or how businesses apply. One is the deep learning bucket and transfer learning. The other one is more traditional machine learning with the boosting and you said, "What's an SVM," earlier. Random [inaudible 00:30:09] things, right? Where it's very important that each client has domain specific information.

Spencer Lourens:

For instance, at CLA we have a coaching program. We have Associates, Seniors, Managers, Directors. We have specific ways in which we organize our folks and every company is going to be slightly different in that manner. So one of the things that we do, this example that's on the slide right now, is showing which one of your clients are at risk of leaving and what we can quantify for your entire client base over a period of time are the most important factors you can see in the top right that were driving that turnover.

Spencer Lourens:

So for instance, this is for a telecom company. So total day minutes, somebody using your service more or less had a very important impact on whether they churned or not. In this specific view, it's not telling me whether it increased or decreased the risk. It's just telling me it's very important. So again, looking at the top right plot here. It's showing me that was the biggest predictor on average, on whether somebody stayed. I would hypothesize that if somebody is using your service more, you're more important to them, on average, they probably don't want to change it because they're like, "Oh, now I've got to go through the pain of finding somebody else. Man, I use this so much. I don't even want to go through that because I'm busy and I just need to use this." The second one is, do they have an international plan? How many voicemail messages they had, was the next one. The number of customer service calls, that's an interesting one, where it may be the inverse.

Spencer Lourens:

Maybe more customer service calls means they have more issues, they're more unhappy depending on how good your customer service department is. Maybe your customer service department's awesome and some people are like, "Wow, every time I have an issue, I call the customer service department, 10 minutes later, I've maybe got a discount from my inconvenience and I've also got the problem solved. These guys are great."

Troy Hollings:

So, Spencer, if I can just pull that thread. So I think that's interesting, but let me make sure I'm understanding.

Spencer Lourens:

Absolutely.

Troy Hollings:

So it sounds like, and I'm just going to throw out simple round numbers here, it sounds like maybe there's, I don't know, there's 25 factors that you could get the data of that relate to customer behavior. And there's almost two sides of the coin. So in my mind, I'm a fan of Netflix and so I think about, okay, if I'm Netflix, there's probably those super users that they're averaging two hours a night watching shows, they're really engaging, clicking the thumbs up button and might be instructive if instead of going from two hours per night for a five week period, now they're at 10 minutes or 30 minutes or they halfway start shows. You're basically taking all of those relevant data points, but then what currently is kind of just gut feel of, oh, well, when we see a lot of watch time, we know it's good, it really... It digs into, okay, well it's multivariate. It's not just those gut feels, it's actually these seven things weighted together that only an algorithm could really identify. Is that a fair Indiana summary?

Spencer Lourens:

I love it. Yeah. And I think there's two pieces to it basically. There's the things we've hypothesized already that our gut tells us are correct, and then confirming them. And then there's things that we may not have thought of. We've at least seen that many times where our client are like, oh, well, that's interesting. Okay. We can change our strategy and how we attack students or employees, et cetera. And that's happened both internally at CLA and with clients. So yes, it's kind of two things that you're attacking, but other cool things on this specific dashboard, because I've looked at one thing out of all the pieces that are here, is what we like to do is quantify the impact. So we can actually tell you, based on your historical churn rate, for instance, we predict your loss from churn could be $357,000. Whereas your outreach costs for an intervention may be 28,000, but you're going to save 170,000 clients through that process.

Spencer Lourens:

So that's a value right there of this dashboard over maybe only a six month period, depending on the retention cycle that we're looking at of 142K that's helping you keep clients engaged with you, maybe even helping you sell more services to your client, i.e cross selling by keeping them in the door. So this is really more of a retention tool, not really a cross-selling tool, but the point is we can quantify through the machine learning, what are the signals that are driving clients leaving, clients staying? How can we help you to develop a targeted intervention plan for reaching out to clients? And I want to get into explainable AI of course, because we talk about explainable AI a lot when we're doing this kind of thing, so that we can know why is a client at risk for leaving us. So I'll stop there because you may have some comments you want to make, Troy, but then I want to talk just a minute about explainable AI and what we would provide in addition to this dashboard.

Troy Hollings:

Sure. I think the only question... So long, long ago I worked at Angie's List, first job out of college, tricked you it's a call center job, so it was a great character building experience. But there was a retention team and so their whole role was to take cancellation calls, but then if they had time to reach out and try to retain the customers, but it was really ad hoc, it was gut feel. So do you see sometimes where, I mean, clients, customers that they actually have a team like a retention team where they have some infrastructure and this is really more just almost using what they currently have, but just calibrating their judgment better or is it always an expense like that 28,000?

Spencer Lourens:

Absolutely. So your point is they've already got resources being directed towards that. So there may be no additional cost, it's just a re-divergent essentially, right?

Troy Hollings:

Right, because I don't know how often just, okay, well here's a hundred people, call them and see how they're doing. But 50 of those calls, maybe it was wasted.

Spencer Lourens:

Right. It doesn't necessarily mean that discounts are involved for that outreach cost, right? To increase that cost, it may just be resources. I agree a hundred percent. And another thing that we've seen where people have a retention team already, they might say, oh, this is being done. And in some cases that could actually be true but what we can provide is a linkage in real time when somebody calls to factors like I said, there's those things we've hypothesized that we think matter, maybe at an aggregate level as what's been looked at before, but now we can target it to an individual customer of what's actually driving that. So you could get that intelligence in real time to know, oh, but this person is actually in this group. They don't really follow the standard rules for why somebody's going to leave.

Troy Hollings:

Got it. No, it makes sense. And then I promised... You want to talk about explainable AI? So, go.

Spencer Lourens:

Well I think I kind of just did it. I was trying to get it in any way.

Troy Hollings:

Okay. Great.

Spencer Lourens:

So yeah, just the idea that... Yeah. Okay. Globally, we told you what was driving retention in that top right plot. A lot of people say that's fine, but what about the individual? When I was in the healthcare industry, everybody wants precision medicine. I don't want the pill or the treatment plan that works for all those people. They're not me.

Troy Hollings:

Yeah. I'm not average. I'm special. I got it. Yeah.

Spencer Lourens:

And we see that, right? We really see that in explainable AI. So the additional thing that we would layer on to this dashboard and that's what we do is for each individual or each business or each whatever you're studying, that's even you could be looking at fraudulent transactions and behaviors that might lead to fraud, we can explain the individual level prediction. Why did the model think that? And it may not have anything to do with the total day minutes. The international plan is yes, because maybe more minutes is good. That leads to a lower risk of churn. It's going to key in on the things that were different about this individual customer that made them more likely to churn, more likely to leave. So I just wanted to point that out, an additional nugget that we try to provide.

Troy Hollings:

I think that, yeah, that makes sense. And I think one of the things, so again, I mean I'm in Higher Ed a lot, and so I've worked on a couple of these projects that your team's built that early warning indicator model to try to predict students that are likely to leave. And I think as we've been having conversations with schools, we've found that sometimes schools have a model, but it's really important... So I'm not a golfer, I'm an Indiana, there's no golfer's in Indiana but I'm a bow hunter. But the quickest way to just do horrible is to lose confidence in your bow like, oh, it's the bow, but it's you and you just devolve into being horrible.

Troy Hollings:

And I think we see that same thing where when people lose confidence in the data, then they stop using the model. And so sometimes it almost feels like it's a little bit of a black box. And I know I talked to one school where it was red, yellow, green, which students most likely, and the red was, well, they're on the football team, they have a 3.0 GPA, I don't think this model's good. But there was none of that further context. So I do think that's probably helpful so you can drill into that specific student and say, what's going on? Why is this specific student having issues?

Spencer Lourens:

Okay, thank you for that. I wonder, I mean, I'm looking at time at 01:40, we've got questions, we have a Q&A session, we've got other use cases. Do you want to answer some questions, go through them and have some fun? Do you want to continue on? What do you want to do?

Troy Hollings:

So we were joking with our nice marketing folks, like, oh, tricked you, it's a two hour webinar. Well, Spencer and I won't have any problem talking so great, but it's still one hour. I think we'll take some questions. You want to moderate that, Spence, and pick out some of your favorite questions and then we'll chat through it?

Spencer Lourens:

Absolutely. And just know that if we do not get to your question right now, we will get an answer out, that's going to happen. So even if we don't do it live right at this moment just for the sake of time, it's going to get answered because there are a bunch of amazing questions and there are some that you guys got me stumped. We should have given questions to this audience first so you could write the ones to stump me. So that's awesome to see. So one that's interesting to me because one of my friends is an FSA in Actuarial Science is from Thurston Group, to what extent has AI and bots been used in the Actuarial Science industry? That is an absolutely awesome question. I know that for the Actuarial exams, they're starting to utilize R as study materials, which is a programming language that often uses machine learning methodologies, but I'm not sure how much AI and bots is used in that industry. I imagine there's a lot of opportunity with all the documents that get mined, but I'm not sure.

Troy Hollings:

I got to call out, I think, my favorite question. Cynthia, "In the middle of the night, when I want to check my investments, how does it read my face in the dark? I know I should be sleeping at that time." Yeah. I don't know. That sounds like a multi-layered problem there.

Spencer Lourens:

Yes. I love that too. I do not know that one either. You stumped me as well.

Troy Hollings:

Yeah. We stumped you eventually. One to one.

Spencer Lourens:

I imagine maybe a night vision type setting. If you train a model to... If there was night vision capabilities on the smartphone to be able to read that, if you could train it to pick up faces in that environment, it's really all about the training data matching the environment of your test data. That's really important. Maybe that would help, but I'm really throwing it at the wall here, Cynthia. So don't take my word for it there. Let's see another one. Oh, what is the difference between deep learning and deep reinforcement learning? That's a really good question from Opt ML. And the idea of reinforcement learning, I'll just say that that's what's been used for trying to learn to play tar games, self-driving, a lot of those types of things where basically you embed typically what you'd have, if you're not doing reinforcement learning is a really large "labeled data set."

Spencer Lourens:

So I have a whole bunch of transactions, some of them are listed as fraudulent, some of them are listed as not fraudulent and I train an algorithm to figure out what's fraudulent and what is not. In reinforcement learning, I have an agent sitting in some environment that's often simulated and I have the agent make an action, so maybe it's turn left and accelerate, turn right and break. And I determine taking that action, did it help or did it make things worse? And through that, I can actually train the algorithm. So they're just different settings to apply deep learning. So that's really the difference. Let's see here.

Troy Hollings:

Can I ask a question, Spence? I'm just being selfish because when you were talking about that chess example, so my understanding of the reason that that was able to happen is because ultimately though chess is incredibly complex, it's solvable. There's an optimal move and if there's not an optimal move at one point, there's an optimal move over three moves, over five moves and there's... It's really hard to run all of those massive calculations, but I've read a book, Thinking in Bets by Annie Duke and she says that life is not chess, which is solvable. Life is poker and her old thesis is that when you're playing poker, you're weighing those expected value calculations, it's like, should I invest in this business? Should I buy this warehouse? Do you think that AI could beat professional poker players on the average? Because that is more nebulous compared to chess, which is solvable. What do you think?

Spencer Lourens:

I don't know.

Troy Hollings:

Two to one, we're winning. We're done now.

Spencer Lourens:

I just want to make a comment because I think this is incredibly interesting. I'm very into AI, right? And so I spent the money to get full self-driving in a car that could drive me. And so it drives me from Omaha to Denver. It drove me from Kansas City to Omaha when I picked it up. It's not perfect though. It makes mistakes. This is why I'm so big on humans as reviewers when we're are applying AI. But what happened with an even more complex game than chess, which is Go, which the like... Action space they call it, right? You've probably heard of this.

Troy Hollings:

Yep. I've heard of that.

Spencer Lourens:

They're alpha goal... Go ahead.

Troy Hollings:

And just to clarify for the audience, that is... I think it's from China and it's a similar but more complicated game, very popular that they were saying could AI beat a human in Go, correct?

Spencer Lourens:

Right. Right. And it did. It learned that. So what I'm saying, I don't want to spend too much time on it because I know we want to talk about the business operations piece, but these models are getting so big in terms of the capacity, that they're able to do really unbelievable things. So I just don't know. I don't know where it's going to go. But I think that leads me into another good question that I found, it was about Bill.com and it's from Crida Services. I think I'm saying that correctly. Nancy Burnett, is it common for AI to not always pick up the correct information? I'm working with Bill.com, and we work with Bill.com too, I think you're talking about the Intelligent Virtual Assistant, which is really awesome by the way. And says, sometimes it gets a vendor name and sometimes not, same with invoice numbers, sometimes it picks up the invoice number and sometimes the purchase order numbers, also on another vendor, it tends to not always find the invoice total.

Spencer Lourens:

And this is fantastic. So we talked about the human being a reviewer. I talked about transfer learning, I didn't talk about fine tuning, but I know I bother you with the idea of fine tuning when we talk fairly often, where I bring that up, right? So one of the things with this extraction of information from an invoice, I showed you an invoice where it basically did it perfectly, but I don't believe it would always be perfect. What we need to weigh is, what is the improvement in time and the improvement in production from that versus the cost of using the AI tool? Because it's not going to be a hundred percent, it's probably going to be 90 to 95% is my guess. So it's our job when we help you deploy something like this, which is going to be very targeted i.e documents come in and we want to extract, knowing the information that you need to use, we're going to make sure that we provide you a way to modify if we get it wrong.

Spencer Lourens:

And so that platform needs to be able to take that information in and if we don't get what you're looking for, Nancy, you can modify it, which I know Bill.com does that. So you already have that. The other thing though, that we would do, that's not always as easily available and it's been about a year or so since I've looked at IVA, again, awesome tool. We can fine tune the model i.e you give us the use cases that are problematic and we can retrain part of that model to make it be able to pick up better in the future. So, that's called fine tuning in the transfer lending literature. So again, from a business operation standpoint, you're like, Hey, it doesn't always get it right that is fairly common. But we have strategies to try to alleviate that.

Troy Hollings:

So spent at the risk of going in the weeds, so I'll pull us out if we get too in the weeds, Indiana remember, but how do we know that you don't... You set up the model and it's really hard to look at a bunch of data and know if it's right or not. So is there an... You do you build a concurrent, like a parallel model that is almost reconciling and confirming? Or do you have... How do you know that you're wrong before you just have a bunch of data and then you wake up one day and you found out, oops, yeah, we automated it, but turns out completely unhelpful?

Spencer Lourens:

Yeah. That's fantastic. So I'm going to loop this into another question on college student retention, because I think we can answer both of the same time.

Troy Hollings:

Okay.

Spencer Lourens:

So, that came from Sarah Lawrence College. Really cool that Sarah Lawrence is on here, by the way. But we would typically do what's called train test i.e we would take a part of the data that's provided to us and we would train our model on it. And we take a whole other piece of your data that has labels so we know like, did a student stay, did a student leave and not use it at all for building the model. Why wouldn't we do such a thing? Why wouldn't we want to use all the data we could to train the model? Well, these models we build can be so good that they can memorize things. Things that don't actually happen in the future.

Spencer Lourens:

They just can fit your data perfectly, but they basically fit the noise because there's randomness. Just because a person has 12 credit hours doesn't mean they're going to leave. Some people take 12 credit hours their entire time throughout college, you have to look at their GPA as well. There's things like that, right? So what I'm trying to say is, when we... You asked how do we know that it's going to work, right? How do we assure that? In one setting, especially when it comes to student retention, those kinds of models, we actually calculate the test performance. We train a model, we hold out data to test it. Data that the model has never seen that has the outcomes in it, so we can only report that. That also allows my team to look, Hey, when I train the data, does it look like it's way more accurate or way better than when I put it on the test data? Then I've overfit my model, is what... That's a concept, it's very well known, called overfitting.

Spencer Lourens:

So I actually need to fit a little bit less complex of a model perhaps, so that it doesn't do quite so well on the training data, because that was false. It was actually overly optimistic. I wasn't being transparent to my person I'm working with in trying to help solve a problem. So that it's more in line how it performs on the test data to the training data. So that's what we would do. So answering that question-

Troy Hollings:

Is that an ongoing process? So if we look at student retention, so that answer made a bunch of sense where there's that first part where you train it, test it, you do a lot of that upfront work so theoretically the model when you deploy it, is good. But I would imagine as we saw with COVID happening or just like I was thinking as you were saying that last week I drove four hours, talked for an hour in Ohio, drove four hours back. If you were taking a weekly average of my drive time, be totally skewed. So outliers happen too. So how do we... Is that a continual training, updating process? You build the model and then need to kind of keep it as time and things change, you need to keep it correct?

Spencer Lourens:

Yeah. You do, because you imagine that the culture changes. If I fit an employee turnover model and leadership really does a good job of implementing the findings, they're probably going to fix some of the systemic things that were leading to turnover. So my old model makes no sense because the issues that were in the data that it was trained on, aren't there anymore. So that's like, covariate drift, essentially is what we might call that. There's a name for that. So there's a whole field called machine learning operations. And we actually need to have some cadence on which we retrain the model.

Spencer Lourens:

In the example about invoices, that would be happening sometimes maybe in an ongoing basis. We just want to continue to try to get that updated as often as we can, because we want to continue to push that envelope. When it comes to student retention, it's often on a semesterly basis that we would look at that. If it's client turnover, that's probably a client by client that we want to have that discussion, but I could see that cycle being much shorter, maybe a month, two months or something like that. And that would just be part of the discussion and set up. Updates at CLA happen anywhere from a year to a month typically on the models that we run internally, for our own purposes, right? So, yeah. Great question. Great question.

Troy Hollings:

Yeah. Makes sense. I think we'll do one or two more questions, I'd also like to point out Michael Swanson says my jokes are entertaining, let's get that guy a gift card, but Spence, you pick one more and then if we have time, we'll do one more after that.

Spencer Lourens:

Okay. Return on investment, another good question from Opt ML. So it really varies right up the business, et cetera, but an improvement of 1% in retention, we've had estimates from certain colleges that we worked with, could be worth as much as like 200 grand. So if we can make a model that helps them target students that we're going to leave, had a high risk of leaving, they don't, that could have that kind of impact. So if you're thinking you're paying substantially less than 200 grand per year for that product and that product continues to update like we were talking about, that could be a return of investment... There's a wide range there, but it could be even upwards of-

Troy Hollings:

Let me pull that thread for a second. So I think it's less of a fixed number and more of a... Almost a percentage of how successful it is. And so you pay a fixed amount, you get a software, let's say it's $30,000, whatever. But then ultimately the ROI is very dependent on how successful it was. If we look at that student example, let's say your tuition's 30,000. So now you save one student, maybe break even, obviously there's scholarships and et cetera, et cetera. But I think we've found that for every... I think this is right. We'll say disclaimer. But for every half a percent increase in retention, a school that we worked with had an extra $200,000 saved. And so I think if you look at that over lots and lots of students, that's where that ROI comes in, but it's a little hard because ultimately it's kind of dependent on how well it performs, but obviously the goal is perform really well.

Spencer Lourens:

Right. Okay. Can I answer two more, very quickly?

Troy Hollings:

Two more, very quickly. Yep.

Spencer Lourens:

There was a great question about unintended bias and AI machine learning. And we need to take that very, very seriously from BECU. Essentially, there's tools and more and more frameworks that are being developed for us to look at that. Are we unintentionally, through an algorithm, treating disenfranchised or minority populations, et cetera, unfairly? Are we just continuing to perpetuate that? So we need to assess those. Those are things we should be thinking about. And then also how effective is Power Automate and some of the products that we may have as part of our office subscription at performing AI? Believe it or not, this is Power BI-embedded, the slide that we're showing right now, is where this came from. So this was actually done through maybe not your current office subscription, but through a Microsoft subscription for not as much of an additional cost as you might expect. You didn't have to spin up this huge Azure Infrastructure and a whole new data warehouse, et cetera. It's using Power BI-embedded. So that's all.

Troy Hollings:

Got it. No, I think that's good. I think we did promise the marketing department we would talk through some of this. So very quickly, I think if we're moving into how CLA might be able to help, I think a lot of times it's really hard to know all these different things. We talked a little bit about data analytics, but there's a whole nother side where it's not even the data science, but it's just getting the data out of your systems that right now is just a horrible pain to get and just getting it in front of your eyes. So as you're going about your day, you just see your KPIs. So there's a lot of different ways that data can help in a lot of ways to wrap your arms around data. And so I think we've found that a lot of times just doing a quick two to three week data assessment where we'd come in, look at all your different data sources, look at what you're trying to do and then ultimately come back and say, Hey, if we were you, these are the things we would probably look at.

Troy Hollings:

And we actually didn't get a chance because we're just going off the top of the head right here. But Spence has a good slide, I'm going to give the 30 seconds summary here. But I think this is the life cycle of an organization as it matures using data. And so at the beginning everything's disconnected. Manual process, no dashboards, everything's all over the place. In higher Ed, I see a lot of schools that they've got Jenzabar, their student information system, they've got Salesforce that admissions is using. They've got Razor's Edge, they've got lots and lots of things. I mean, it's honestly the first step, if your house's on fire, you got to put it out.

Troy Hollings:

Phase two is kind of integrating data a little bit more, taking it to the next level, getting less manual process, maybe automating a little bit of those data flows. Phase three and phase four though, that's where that data science starts to get wrapped in. But I think the reason that Spence... He called me and just yelled at me about this in a happy, good way but basically I think if we don't think about the data science from the beginning, you get to the end and then you're like, "Well, maybe we need some data science here." But I think that all of that to say it's complicated, lot of times we'll start with an assessment and I'm hoping that we don't bite off more than we can chew here, but got it cleared. And we're offering, for anybody who wants to, a 2 to 3 hour workshop session with someone from our team.

Troy Hollings:

So basically no charge. Free of cost. We would just kind of come in and either in person if it makes sense, virtual, if it makes sense, and just kind of chat through what are all the things related to data that you all are trying to accomplish. So I think somehow somewhere in this magical go to webinar thing, there should be a button that pops up and says I want to be contacted. So that sounds like something you'd be interested in? Click it. We'd be happy to reach out. And even if you just want to say hello, we'd love to talk to you. So any final closing thoughts, Spence?

Spencer Lourens:

No, I just want to thank everybody for coming.

Troy Hollings:

Yep. Thank you all. If you have any other questions, throw them in the question spot, we can reach out over email and answer them. We love to do this stuff. So I really appreciate the attendance and thank you very much.