LEND ACADEMY

PODCAST TRANSCRIPTION SESSION NO. 116-SCOTT ZOLDI

Welcome to the Lend Academy podcast, Episode No. 116. This is your host, Peter Renton, Co-Founder of LendIt and Founder of Lend Academy.

(music)

Today's episode is sponsored by Wunder Capital. Invest in commercial scale solar energy projects across the US, earn up to 8.5% annually while also diversifying your portfolio, curbing pollution and combating climate change. You can begin investing with as little as $1,000 and best of all, Wunder Capital doesn't charge any investor fees. To learn more, create an account for free at wundercapital.com/lendacadmy. Wunder capital, do well and do good.

Peter Renton: We have a fascinating guest on today's show. I'm delighted to welcome Scott Zoldi, he is the Chief Analytics Officer at FICO. He's also got a Ph.D. in Theoretical Physics which at first glance might seem a little strange, but as we discuss in this episode you'll see how it fits actually perfectly because Scott has been analyzing data in complex systems his entire career.

So in this episode we talk a lot about artificial intelligence, this is something that FICO has been doing for many decades, Scott has been doing for many decades and we really get into how FICO is using artificial intelligence today. We talk about the many patents that they've received, we talk about fraud and how their system of fraud has been so successful, we talk about alternative data and the different kinds of data they're using and we also talk about Scott's view on where the FICO score is at today compared with ten years ago. It really was a fascinating episode and I hope you enjoy the show!

Welcome to the podcast, Scott.

Scott Zoldi: Thank you very much, Peter.

Peter: Okay, so I like to get these shows started with giving the listeners a little bit of background about yourself. I know you've been at FICO for a while, but can you just give the listeners just a background about yourself, how your career has gone?

Scott: Sure, so I have been with FICO for 18 years, currently, I am in the role of the Chief Analytics Officer. My training has been in theoretical physics so prior to coming to FICO, I was studying chaotic systems, turbulence, doing work with the government labs and I decided to pivot into analytics around financial fraud and credit risk when I became aware of the amount of data that was available and some of the really interesting problems that we have from a machine learning and analytics perspective and came to FICO in 1999 and have been a scientist ever since going through the different ranks from an entry level scientist to today being the Chief Analytics Officer at FICO.

Peter: Huh, so you've really been looking at data and sort of non-linear type data your whole career, it sounds like.

Scott: That is exactly right and, you know, when I was in my Ph.D., I was looking for phase transitions between let's say fluid systems and now one of the areas I spent a lot of time on was actually fraud detection and you can imagine that being some sort of phase transition between normal and abnormal behavior. I saw those parallels right away, but very often it sounds a little odd to have a physicist, a theoretical physicist in this business.

Peter: Oh, the way you explain it right there makes perfect sense to me so that's really interesting. So I just want to talk a little bit about FICO, the company. I mean, everyone has pretty much heard of the FICO score. Not everyone knows exactly what you guys do so can you just explain what FICO actually does, what's your business model, what products and services you offer?

Scott: Yeah, so FICO is a company which is an analytics company, we've been around for more than 60 years now and our entire business is based on the development of models that allow businesses to make more intelligent decisions. Those decisions can range from making a decision around the probability of fraud to measuring a credit risk or a default rate and making a prediction there to things like optimization of systems, marketing, cyber security, what have you.

It's a company that's constantly growing, we spend a lot of time on the development of new analytic methodologies in machine learning and other areas because we really want to kind of optimize and improve our analytics over time for the business problems that we work on for our clients.

Peter: Okay, and then so how does FICO kind of interact with the credit bureaus because obviously everyone knows the three big bureaus here, they each have their own score which some people call the FICO score. Can you just explain the relationship that you have with the three bureaus?

Scott: Yeah, so for each of the three bureaus the relationship consists of the fact that these bureaus are collecting the credit histories associated with consumers and the payments that they make, any delinquencies that get recorded associated with these consumers and they come to FICO for the development of the FICO score, this is the score which is used for making a decision around delinquency and credit risk. Each of these bureaus have a credit score, FICO score that is tuned based on the data assets they have. Sometimes they might have slightly different data based on the reporting of the credit history to each of the different bureaus. They provide us data for the purposes of building our models and then customizing it based on the data that they have at each of these bureaus.

Peter: So you're actually selling them your credit model or your analytics model, I guess, that they then...or are you selling them the output, like the score or...what are you actually providing?

Scott: Right, so we're providing the score so it works on the data they collect and then we're providing the model which produces the score that they can then provide to their clientele.

Peter: Right, okay, okay, got it. So you mentioned that you're Chief Analytics Officer of FICO and it's not a title that you see all that often. Can you explain what your role entails?

Scott: Yeah, as Chief Analytics Officer I have responsibility for the kind of analytic roadmap approaches to our business and products. So a lot of that entails having a large number of Ph.D. data scientists reporting to me that are responsible for the development of the models that FICO produces.

Other parts of it consist of a research plan related to new algorithms that we're developing, filing patents around some of the inventions we create in the analytics space and then also focusing very much around anticipating what our clients may need in the future from an analytics capabilities perspective so really trying to make sure that FICO is at the cutting edge of the analytics practice, that we continue the legacy that we've had over the last 60 years of using empirical models to make better decisions for our clients.

Peter: So you mentioned patents there, I'm curious about that. Can you give us some examples of some of the patents that you have been awarded?

Scott: Yeah, so it's interesting, we have patents, more than 130 at this point that have been granted. You know, some of the ones more recently that have been granted are focused around unique ways of allowing models to self-learn in production. So it's a class of models, for example, that are really important in areas like cyber security or fraud detection where you may not be able to...let's say in fraud detection you can build a really great model because you usually know where the frauds are and who is not fraud, you can build let's say a traditional machine learning model, but as we all well know, fraud is a very kind of rapidly changing problem and so these models have to adjust over time.

So we have these technologies called Adaptive Analytics that ride on top of our machine learning models and then learn what is happening in the production environment. We also have another set of patents related to completely unsupervised models so models that you wouldn't necessarily have CAGS, let's say good and bad or fraud or not fraud or cyber breach or not cyber breech and the models have to learn based on the data that they are presented what normal looks like, what abnormal looks like based on the data that it processes over time so it can kind of serve up examples of let's say customer behavior or machine behavior that might be risky so that someone can go investigate that.

With more and more of the data out there streaming at us at incredible rates, most of these technologies that we develop from a patent perspective focus on these real-time analytics systems that can process transactions very, very rapidly, but produce these analytic scores even if they're self-learned in production.

Peter: Okay, so that brings up an interesting point because I went back and watched your panel, there was a panel at LendIt that you were on that focused around AI and you said this a couple of times in the panel that you've got to be able to justify to regulators why you are rejecting somebody on a loan application or why the score is what it is. So I'm just curious about...when you're talking about self-learning models that are sitting on top of your regular machine learning models...I mean, isn't that somewhat opaque and difficult to explain?

Scott: It really depends on the architecture of the model that we develop. In general, any machine learning model today it is usually difficult to explain because they're block boxes, but when you have a model that's adjusting over time, so for example, if I go to one of these adaptive models for fraud, the model itself, the adaptive model, essentially would be reflecting on the most recent fraud and not fraud that would be seen in production.

In terms of explaining that output, we can develop the model in such a way that it might be much more like a scorecard model that's learning continuously based on feedback, and those are interpretable models. So, in general yeah, it is tricky, but if you're maintaining the history on which you build that model, even an adaptive self-learning model and you choose an architecture which is explainable, then you don't run into the same sort of issues from an explanation perspective.

Peter: Right, right. So I'm just curious, this question came to me when I was watching your panel. Would the model be much better if you didn't have to be able to explain it or you think it really is just as good knowing that you've got to be able to explain the decision?

Scott: It's a great question. You know, in terms of...I'll answer it this way, it really depends how well you understand the problem. So, for example, the FICO scores we have in the US, right, we've been building these for 25 years, we've probably explored almost every perturbation of the way that we could define features and variables for that problem. So what we find in those areas is if you try to apply machine learning on top of this model that's been refined over 2.5 decades, it doesn't find a lot more. It's not warranted to take on the expense of let's say having a machine learning model that provides very little lift and improvement and then introduces complexity that you have to try to resolve with the regulators from a machine learning perspective so that would be one use case.

Another use case could be a score that's being built in a different area of the globe where there might be different data sources, less domain experience and that might be a situation where the machine learning models would be used to gain more insight because they wouldn't be able to learn the difference between features and speed up that learning process. We'd still have to defer to the local regulation whether or not you use machine learning to gain insight and build better models like learn features and segmentation and still use an interpretable model like a scorecard. But that's where the machine learning could have some benefit if you're exploring entirely different data sets or types of data or a set of clientele that you haven't been working on for decades but it is a brand new area and there might be different relationships to learn there.

Peter: Right, understood. So still on modeling here, it sounds like just from the way you describe your background...I mean, that FICO has been doing machine learning for decades, it sounds like, because obviously regression analysis, linear type regression analysis, has been standard for decades and then we've moved it seems fairly recently from my perspective…but obviously, I'm not as close to it as you are, so have you been using this sort of artificial intelligence sort of mode of analysis for a long time? Can you explain when you moved over to that or has that just always been the way you've done it?

Scott: So it's been something that's been at the heart of our company for a very long time. Our fraud solution, Falcon, it's a neutral network-based model that monitors about 2/3 of the world's payment cards for fraud...that was introduced in 1992, and that's when we brought these machine learning models, neutral network-based models along with the behavioral analytic profiles to the banking industry to kind of help fight an ever growing sort of problem that was really a tremendous challenge for the banks in the very early 90's.

So we're celebrating our 25th anniversary of not only using AI and machine learning, but productizing it and today, models like Falcon are almost considered part of the infrastructure with respect to a trusted sort of methodology for maintaining fraud risk and detecting fraud while minimizing false positives. But our history goes even before that. We were actually using machine learning models in the defense area in the 80's where we would go and do some of this work trying to detect let's say what would be a tank in a landscape and those are very basic sort of image recognition problems that were used from a machine learning perspective.

Today, we have probably....well we have more than 70 patents just in the machine learning AI area from that legacy so it's always been part of our pedigree. And even in areas like credit risk where scorecard technologies are very palatable and proven technologies that are transparent and explainable to regulators, we still use those AI techniques to explore different features and segmentations. So we consider them a part of our tool sets and parts of our business we deliver the machine learning model as an output and in other parts of our business we use machine learning in the development of our models, but we deploy them in a different way.

Peter: Right, okay. What do you think…it seems like the last two years, maybe three years and even more so, I feel like even in the last six to nine months, that artificial intelligence is all the rage, it's everywhere, you hear about it, there's articles being written daily about...in financial services it seems like everyone now has this advanced AI that they're implementing so do you sort of sit back and laugh when you sort of read some of these articles? What are your thoughts on it?

Scott: So one is I'm excited that machine learning and AI is the rage again so I think that's a really positive thing because it's a tremendously interesting set of technology. You know, some of the claims out there are sometimes a little bit laughable because I think a lot of these new companies that are popping up are not aware of some the legacies like in the fraud space that banks have been using these neural network models for two and a half decades. But, you know, I generally sit there and I pause because what I worry about is it takes a lot of care to properly productionize and operationalize AI. So from making sure that models aren't overtrained, right, to making sure that you have the proper amount and degrees of freedom, that you really thought about the data you bring in to the machine learning models so it doesn't learn relationships that might cause bias or might cause the model to not perform in predictable, reliable ways.

So that's what I actually spend time worrying about because at FICO we have a lot of experience with this and we treat it with a lot of care and we have a lot of governance around these models in terms of how to develop, we encourage our customers to apply governance once these models are installed. My only concern right now is that with this hype cycle where AI and machine learning have to be part of everything that we interact with, that they be built properly. Part of that, frankly, is with big data. A lot of people look to machine learning as being kind of a savior for all these investments they've made collecting data over these years and, you know, one of the worst things we can do is just blindly apply these technologies.

So whenever I'm out there and I hear about different companies working this space, I just always encourage them to really take some time to think about the problem they're solving, the data they send into the AI, machine learning and have some governance around it. I think we'll see a lot of really great applications, without that, we'll see a few models go sideways for different folks or different startups. But I think that will just become more mature over time as some people start to learn some lessons with respect to the AI that they're using and we talk less about algorithms and more about proper use and governance of the technologies.

Peter: Yes, that makes sense. Let's talk about that data then for a second. There's been an explosion in data in recent years and like you do need obviously pretty complex systems to be able to make use of it all. I'm curious to get your take on it, I mean, so what data are you using? The second part of this question is what's considered alternative data...you can't see me, but I've got my fingers in quotes like rabbit ears so alternative data, social media, smartphone data, even alternative financial data like rent, mobile phone payments, utility payments, all that sort of thing. So, have you greatly expanded the data you bring in to your models?

Scott: We're most definitely looking at increasing the types and ranges of data that are used. In fact, we had a FICO financial inclusion initiative focused on development of new scoring products, partnerships and services to really kind of make sure that we can bring in let's say more of the 3 billion customers that are underbanked globally or unbanked. So we feel this is very important because at the core of let's say if we stick to a FICO credit score, the core of that is bureau information. Even though the FICO score is in over 25 countries right now, the quality of bureau data and the availability of data in different parts of the world are really different. So part of this is to explore what's available from a data perspective to improve this score over time.

As an example right, we had an initiative which was kicked off in 2014 around financial inclusion and that's where we have this FICO XD score, where we said for those that are unbanked or underbanked in the US, we would go leverage things like rental information, whether one is paying phone bills because those are pretty good proxies for let's say paying off a line of credit also and we found that that has allowed us to extend FICO scores, these XD scores, that allow people to kind of enter that kind of credit cycle and build up a credit history. Now with that said, that's a special use case, we look at the core business, we also have to look at what regulations say about that.

For example, today we're not necessarily using mobile device information or social media primarily because of regulatory concerns around the proper use of this data with respect to a traditional credit data application, particularly if you put it into a model and you're not entirely sure how. Let's say my last tweet has an impact on how my credit score will go up or down, and so we explore it, we explore it all the time, but it always has to do with a view to what are the regulations that we're able to use here…would the model, would we be in line with regulation in terms of type of data and would we ensure that we're not building models that learn relationships that really aren't going to be helpful in terms of impacting consumers in an unintended and maybe misunderstood way.

Peter: Right, right. So is disparate impact the main concern here where you think you're going to sort of use data that is just by the very nature of the data itself is going to discriminate against a certain group or is there something more?

Scott: So that's one of the really big concerns because particularly if you look at big data or lots of different data sources, the correlations get learned, it could potentially be discriminatory or could be very correlated to let's say some sort of protected class, but just in addition to that, the data that gets collected will change over time, may not be steady at the same level of rigor and the models might just learn things that are really not real relationships.

As an example, if we thought about a fraud example, if we move away from credit risk for a moment, you know, if I collected data…there's a lot of fraud reported in New York and I build this model, the model might assume that any time a transaction occurs in New York that it's fraud and that's not a real relation, that's a correlation of data, but it's not causal.

One has to really focus on those and we want to avoid those sort of things being learned as we combine lots of disparate data. As we look at data, we have to constantly ask those questions of ourselves and ultimately look at some of the relationships that are being learned so that we have a level of confidence that they are legitimate relationships and ones that we could defend and be comfortable with.

Peter: I want to go back to the FICO XD scoring model that you mentioned. Is the output of that model…that you score these people that you don't have the traditional bureau data, you talked about...rent payments, mobile phone payments and stuff like that, but I'm just trying to figure out how that is different to the traditional ones that you'd get at TransUnion, Experian and Equifax.

Scott: The main thing is that is different, you know, our traditional score is leveraging the credit bureau and the data within the credit bureau, but if we have someone that's kind of new to this country or a young person or someone that just hasn't entered into the banking environment, there really would be a very shallow or no credit bureau information whatsoever. So they would be unscorable by the traditional FICO score. And that's where we have a relationship with LexisNexis and with Equifax around collecting other alternate data which now enables us to build that score based on the kind of rental and telco data.

We build it in such a way that it actually has a strong correlation with the FICO score and that's the entry point. Where now someone, instead of having to make a decision without a score, can now leverage a score which would be interpreted on the same scale as a traditional FICO score and they enter, they enter into the credit life cycle.

Over time, they start to develop their own credit bureau history and they would pivot off a FICO XD score to the traditional one because now they started to have a credit bureau and they have a history of paying bills and there's a history being learned on that and they would pivot back to a traditional score at that time.

Peter: Okay, I also just want to follow up on the fraud piece because I think it's really important. I want to just get a sense, you said the Falcon product has become sort of ubiquitous in the banking industry, how has it become so successful, what are the inputs that you take that are so predictive when it comes to fraud?

Scott: Yeah, so what made it very, very successful in the early 90's was back in those days, the banks were primarily writing rules and kind of large sets of rules to detect fraud. Fraudsters were able to very rapidly sort of learn the rule thresholds for when let’s say a case would be generated or not generated so the fraud was spiraling out of control.

The other aspect of this was as the banks were trying to go figure out all of the different ways to stop this fraud they would write more and more rules. In doing that, what you end up doing is maybe not detecting much more fraud, but with each and every rule that you write, you add more false positives. So they got into this situation where they wanted to detect more and more fraud, but false positives were a huge issue so that's where we became very successful because we applied a model which was kind of finding the root, kind of best relationships between lots of different input data that would occur on the offstream and also features that were derived based on a transaction history.

It did two things; first thing is it was a lot harder for fraudsters to try and figure out when a case gets generated because it's a much more complex decision based on these relationships that the model's learning based on historical data of fraud and not fraud versus something that's a little more simplistic in terms of rules. The other aspect beyond fraud is just trying to guess or work their way around it was that the banks were finding very, very high levels of detection and much, much lower levels of false positives.

We were very fortunate because we installed this first at First Data and we had an ability to then allow many, many different banks to leverage this model and as they all saw the benefits of this they decided that one of the best ways to have the most predictive model, from a machine learning perspective and we helped them with this, is for banks to allow the data to be aggregated by FICO for the purposes of building a model that teases many examples of fraud and not fraud so that this model would be better than any model they could build on their own because they see a much broader view of a set of good and bad data over time.

Peter: Right.

Scott: That's really a big part of why it was successful over time. You know, you turn the clock ahead 25 years, it's still very, very successful although many of the banks talk less about detection of fraud. They talk more about improving the customer experience and so it's more about when you do have a false positive, make sure it's a false positive that makes sense, that it's a new behavior versus let's say just a risky behavior.

So the technologies continue to advance over time because we want to continue to keep fraud at an acceptable level and then make sure we're impacting as few people as possible and really learning their individual behaviors very, very deeply.

Peter: Right, that makes sense. We're almost out of time, but there's a couple of more questions I really want to get your perspective on. There was an article I think it was a couple of months ago now that talked about how the financial crisis now, we've gone more than seven years, and the bankruptcies are rolling off and the bad behaviors that happened during the financial crisis are no longer part of people's credit reports so they're saying that FICO scores are now higher than they've ever been before. I read something that there's more people above 600 than ever before so the question I pose, because you've got this definitive number, is a 700 FICO score today the same risk as it was 10 years ago?

Scott: The 700 score today is not the same as it was 10 years ago. If we look at where we were 10 years ago, we were on the cusp of the subprime mortgage meltdown and then the recession. So there are more 700 scores, they are a reflection of an improving economy over the past seven years and that means people have been able to pay their bills, they've been able to manage their debts in a much more effective way and many of them have seen kind of rising equity from home prices.

So that has all kind of made for much better behaviors from a credit risk perspective over time. Even though the risk would be different between a 700 today and a 700 ten years ago let's say, it's still an effective ranking of risk. It's just that there are going to be these factors from time to time that are kind of external, much more kind of global that can cause a shift in terms of the relationship between the FICO score and let's say a probability of default. So that does occur, but that's generally what's driving the fact that we see more 700 scores today is just people's lives are in a much more kind of financially healthier state when it comes to paying their bills and managing their credit.

Peter: Right, okay, that makes sense. So last question, what are some of the things that you're working on personally today that are exciting for you and things that might be coming down the road for FICO?

Scott: One of the areas that I'm really excited about is explainable AI so this is explainable artificial intelligence. So when we look at all this excitement around AI and machine learning, we always go back to the fact that many of these algorithms are not explainable or transparent. FICO is actually...we just had a patent that expired that was written in 1997 around explainable AI, so somewhat of what I'm really excited about is to continue to push the frontier of explainable AI for a couple of reasons.

One is because that would allow more confidence in the use of machine learning in areas that are more heavily regulated such as credit risk, but also if we look at the EU. We have this general data protection regulation that also states that if you're making a decision on a consumer that you need to be able to explain the reasons for those decisions so this is a big research area for FICO.

It always has since we introduced Falcon and it continues to be one that I think if we crack that nut and really come up with very good explainable artificial intelligence then we're going to see a lot more responsible use of this technology over time. So that's one of the topics that I'm really excited about to kind of clear the way for meeting regulations, not just credit risk regulation, but regulation across the globe related to the use of models for decisions.

Peter: That's fascinating, well unfortunately, we have to leave it there. I really appreciate you coming on the show today, Scott.

Scott: It's been a real pleasure, Peter, thank you so much.

Peter: Okay, see you.

You know, sometimes I think FICO gets a bit of a bad rap in the fintech industry. You hear companies talking...oh, we don't really use FICO anymore, we're beyond FICO. I think we can all appreciate, after listening to this interview, that FICO has certainly been on the cutting edge when it comes to artificial intelligence, anyway, for a long time. They were doing artificial intelligence before many of these new CEOs in fintech were even born, I think, it's good to remember.

But I also think I'm excited about the explosion in artificial intelligence and how we can really create a better mouse trap, we can create better models, we can create better predictive models for all kinds of consumers, not just prime or mid-prime or subprime, but all across the spectrum and we can do it internationally. I feel like this is an exciting time, there's going to be a lot of advances coming down I see in the next couple of years and it's going to be the fintech companies that are really doing this responsibly and in an intelligent way that are going to be the winners.

Anyway on that note, I will sign off. I very much appreciate you listening and I'll catch you next time. Bye.

Today's episode was sponsored by Wunder Capital. Invest in commercial scale solar energy projects across the US, and earn up to 8.5% annually. Bryan Birsic, the CEO of Wunder Capital was a guest on our recent podcast, Episode No. 104 and just worth noting, they've financed 40 large scale projects to date with no write-offs whatsoever, so far. To learn more and to create a free account go to www.wundercapital.com/lendacademy.

(closing music)