



PODCAST TRANSCRIPTION SESSION NO. 192 / DOUGLAS MERRILL

Welcome to the Lend Academy Podcast, Episode No. 192. This is your host, Peter Renton, Founder of Lend Academy and Co-Founder of the LendIt Fintech Conference.

(music))

Today's episode is sponsored by LendIt Fintech USA 2019, the world's leading event in financial services innovation. It's going to be happening April 8th through 9th, at Moscone West in San Francisco. We're going to be covering digital banking, blockchain, financial health and of course, online lending, as well as other areas of fintech. There will be over 5,000 attendees, over 250 sponsors and registration is now open. Just go to lendit.com to register.

Peter Renton: Today on the show, I am delighted to welcome Douglas Merrill, he is the CEO and Founder of ZestFinance. Now ZestFinance has been around for a few years now and in that time they have really built a reputation as one of the world leaders when it comes to the development and deployment of AI tools in the credit and underwriting space and their ZAML software, which we talk about, has really created some breakthroughs around explainable AI.

We go into that in some depth, we talk about some of the implementations that they've done, we talk about the data that goes into some of these models which makes it more predictive and we also talk about Douglas' relationship with regulators and much more. It was a fascinating interview, I hope you enjoy the show.

Welcome to the podcast, Douglas!

Douglas Merrill: Thanks so much for having me.

Peter: Okay, so I like to get these things started by giving the listeners a little bit of background about yourself. You have had an interesting career so could you give us some of the highlights.

Douglas: Yeah, I think most people would describe me as having a somewhat random career. (Peter laughs) I've got a PhD from Princeton, Princeton gave me a PhD in AI a long time ago, then I worked at a place called the RAND Corporation which is a think tank in Santa Monica, California where I did a lot of fun work, but notably it's on the beach so you could walk out the back door and go on to the sand to go swimming and I spent a lot of time that I should have been in meetings out that door. (Peter laughs)

Peter: That must have been fun.

Douglas: Yeah, it was great fun, it just maybe wasn't great for my productivity.

Peter: Right.

Douglas: Ultimately, I ended up at Google where I was Chief Information Officer for several years, but probably my most random job in my background was I was the Head of EMI Records,



LEND ACADEMY

which was the fourth largest record company in the world when I was there. It has since gone bankrupt, but it went bankrupt well after I left; I want to make sure that's very clear. (Peter laughs) After EMI, I founded Zest, I've had a fun career doing all kinds of interesting things with all kinds of smart and fun people.

Peter: Then so tell us a little bit about ZestFinance and...maybe start with how you got the idea for the company.

Douglas: I founded Zest in about 2009 or so, with a mission to make fair and transparent credit available to everyone. The operating idea was credit underwriting was a (inaudible) of modern life and yet about 80 million Americans either have no relationship with a bank or a poor relationship with a bank, 800 million Chinese are in the same situation, the numbers are similar for the EU and even more so for Africa, and I was interested in a world where so much has changed in the last couple of years in technology, particularly. How is it the case that underwriting hasn't changed since 1950?

Before 1950, to get a loan you basically would go to a bank and you'd sit across the big mahogany table from a man and they were always men and they were always men wearing blue suits and red ties and you say, hey, I want a loan. And that man would say, oh, you know what, my kids played basketball with your kids, you're a good person, I'll give you a loan. That worked great unless your kids didn't play basketball with his kids in which case you didn't get a loan, that's kind of unfair. Two gentlemen, one named Fair and one named Isaac, whose names later came together as Fair Isaac had a brilliant idea of using logistic regression and credit bureaus to come with a standardized credit score and that just changed the world.

I think it's one of the most important innovations of the 20th century and credit availability went up markedly, losses went down markedly. It created this amazing, amazing positive outcome and I think everyone should thank FICO for the structure of the world today, but FICO has some weaknesses and the primary weakness is in the math itself. Logistic regression works well if you have every piece of data that you need and it's all correct. In the event you have missing data or it's incorrect, logistic regression fails in a kind of unpredictable manner.

I believed when I founded Zest that the primary reason that 80 million Americans were having trouble getting credit is that in many cases, on their credit bureau, the information was either missing, this is sometimes called a thin file problem, or there were errors. And so some number of them were actually good credit, they just didn't have good FICO scores and we could use the math that we built at Google to do a better job of finding that and that turns out to be true.

In the last ten years, our clients have used our software to make material improvements in credit availability. This is not a ding against FICO, I still think FICO is the most important company in credit today, but I also think that lenders have more options because of machine learning.

Peter: Okay, so then...I guess you started obviously doing this. This was before really anyone was really talking about, you know, machine learning and certainly it wasn't top of mind of many



people in the credit space. So how did you kind of start the company? I think I read somewhere that you actually made loans yourself initially so tell us a little bit about those early days.

Douglas: Yeah I mean, I think people who came out of the machine learning companies, you know, places like Google or Amazon, would have recognized my company as another ML company so I feel like I came out with the notion that I wanted to build an ML company in credit. However, I needed to demonstrate to the world that that made sense so the way we did that is for a couple of years we actually made loans using our algorithmic structures and our tools and, you know, they worked pretty well and then we ultimately stopped but the problem with being a lender is it's capital inefficient, it takes a lot of capital and the capital is expensive.

Peter: Right.

Douglas: The reason that banks make loans is they have very cheap capital so I didn't really want to do that long term. I did it long enough to demonstrate our tools and then we ended up serving a couple of clients who themselves made loans and we did that in a fairly bespoke manner for quite a while and then when we moved to our broader product which was the same across the broader market and you can kind of see the trend there as the financial world caught up with, oh, I understand machine learning, I understand how these tools might work, we progressively got more and more towards the SaaS model that I actually wanted to start when I founded Zest.

Peter: Okay, so then, I'm just still curious about those early days, I mean, how were you able to convince other lenders that your solution was better than what they were currently doing?

Douglas: Ultimately, that's just math. So you give me a sample of loans that you've made, whether it was a good loan or a bad loan, and we'll build a model for that. And then you give us a set of loans and you don't flag them, so you don't tell me whether they're good or bad, and I just predict whether they're good or bad and then you go count...

Peter: Right.

Douglas: ...tell me how many the model was right with, how many was it wrong with, ultimately that's just math and it's pretty well understood math. Our tools, we do that computation automatically, we do a bunch of economic parsing automatically so you can say, oh, if we had done this model we would have saved \$100 million in losses, or would have gained \$40 million because of increased approval rates. Those things are just math.

I think harder than that is getting people to sort of think through, wow, you're going to use 500 variables instead of 20, we all understand 20 variables; 500 variables, hey that's a little scary. It just turns out to really help in those cases where you have thin files or you have erroneous data. Having more signals, more variables helps reduce the impact of those incorrect things.

Peter: So then let's talk about that for a second. What are some of the data points that you're using today, or maybe even back then, what were the data points that you used, I mean, maybe



just give me some examples. I know it's not just one data point, there's hundreds, as you say, but what are some of the examples of that that you found really helped lift the performance of the model?

Douglas: So one of the things which is true when you have 500 signals is that really, really good signals, still don't matter all that much because every signal is worth, you know, 1/5 of 1% so a great signal is worth a very small amount. Early on, we discovered that people who put their name in in all uppercase were substantially riskier than people who put their name in correct case. So for example, if I put my name capital D O U G L A S, I'm much higher risk than if I put capital D lowercase... o u g l a s, I just misspelled my own name which is fairly funny (Peter laughs).

So if I fix that, I don't exactly know why that's true. I can make up a story that oh, maybe people who are putting their name in proper are better rule followers but your applying kind of post hoc reasoning to something you don't know why it is, but that's a real signal in a real model. Similarly, one of our recent clients discovered that if you put in the company name that you work for, including the company type so instead of putting in ZestFinance, you put in ZestFinance, Inc., you're a higher risk. So being more correct in the company name increases your risk, well, why is that? Well, you don't really know, but maybe that increases the likelihood that you're a bot because no human uses ", Inc."

Peter: Right, right, yeah, that makes sense and as you say, it's not like someone...obviously there are many, many good people will actually put that in like that, but it's just, as you say, it's one of 500 signals. If the other 499 signals or 498 signals are good, it won't make any difference, right.

Douglas: That's correct because you end up adding them all up so if that one is a little bit wrong, well, alright the others are probably a little bit right and ultimately, it gets zeroed out in some sense.

Peter: So then one of the big things today that is talked about a lot when it comes to AI is this thing about explainability and you know it's something that I know has made many banks and actually many lenders, in general, reticent to embrace this fully because you've got to be able to provide a reason to the borrower why he was declined. So tell us about a little bit about the advances that you guys have made in explainability.

Douglas: So first of all, it's a valid worry because the banks in the US are required to give adverse action notices. If they decline you, they have to tell you here are the top five reasons why we declined you and it turns out that process is not ideal, but it's still required. Also, you're required to periodically demonstrate that you don't have a disparate impact problem where you don't tend to give loans more to men than women or any other of the protected classes.

Doing those things requires understanding what's in your model. Understanding what's inside that logistic regression model is trivial, you just read the coefficients. Understanding what's inside a score picker machine or a massive gradient boosting machine, those are hard so we



really did need some big improvements in understanding how to describe what's in the model. There have been over the last few years, some improvements.

For example, there's some really simple technique called permutation impact where you just bury stuff and see what happens; it's like when you bury one thing and it doesn't matter, it's probably not important. That turns out to be a pretty, you know, gross view of the model, it doesn't have much kind of fine grain and it's pretty hard just to run against big models, but that was a really big step forward where someone realized, oh, let me just pull on some strings and see what happened, good insight.

I think that increasingly the world is demanding a much finer grain understanding of what's in a model because we're increasingly caring about implicit bias and other topics around there where you could, from a potentially regulatory problem and from a moral problem, back into something where you might be disadvantaging groups that you don't want to do and that requires a more fine grained approach. Arguably, a much more fine grained approach, and there's some other things to it as well, but ultimately, you need to be able to use a tool better, you need to look across the table at your financial services client and and say, yes, I guarantee you will understand what's in your model, you'll understand how it works and that will be able to work within a real-time.

Let me give you an example about what I mean by real-time. So there were two amazing basketball players kind of in the 1980's, one named Michael Jordan and one named Scottie Pippin, they played for the Chicago Bulls together. When they played with the Bulls together, they won, I think, six league championships, I think Michael Jordan got three or four MVP's, he was the shooting lead for many years, he was just a star on those teams. So I often like to ask the people I'm speaking to, for a show of hands, who is the better player, Michael Jordan or Scottie Pippin, nobody ever chooses Scottie. That poor guy gets no respect. But if you look at what happens after they've played together it's kind of interesting.

So after they play together, Michael Jordan goes to Washington Wizards and proceeds to have two massively losing seasons where nothing good happens to the poor guy and he retires. Scottie Pippin goes to the Trailblazers where he proceeds to have three seasons where they make it to the finals every year, they don't win, but they make it a long way, Pippin's got an MVP, Pippin has a great career.

Now you've got to ask the question, oh, hang on, who's better, is it Scottie Pippin or Michael Jordan? It made a lot of sense in one context, the answer was obvious when they were together. In the other context when they're apart, the answer is also obvious, but they're different and really thinking through in a complicated machine learning model how to handle that difference is the core of the explainability problem.

Peter: Okay, so I guess say more on that like put it into a credit example. A customer can send an adverse action letter to their borrower declining them, how does that sort of translate?



Douglas: So today when banks send adverse action letters or send these decline notices they more or less take the reasons from one of the credit bureaus and they just sort of package them up into a physical letter and send it off. Some banks do a much more thoughtful job of this, but some banks, if they're a little bit rote or whatever, that credit bureau said, I'm just going to pass on. In machine learning, it's hard to figure out what those things should be, right, you've got 500 signals and they're all relatively unimportant. How do you find the five that are most important?

So that involves using the permutation impact idea, taking 500 signals and randomly permute them against each other which requires something like 500 factorial computations which is a number larger than the number of atoms in the universe. It's a very hard problem to solve that way. Then you look at the ones that bubble to the top that say, oh, if this number had been higher or that number had been lower, your score would have been high enough that we would have granted your loan and then you've got the five approval reasons. The problem with that approach is that you don't have enough time to have more computations than there are atoms in the universe...

Peter: Right:

Douglas: That's just hard and where I come from just hard is the problem you don't want to solve. So there are optimizations, as you can do, that yield both fine-grained, better fine-grained solutions and are more computationally trackable, but that's a silly thing for a problem of how do you understand high level interactions between signals, between variables to know which of the five that are more important.

It always comes back to the same question, but if you look at variables alone, i.e. you look at Scottie and Michael together, you see one picture of the world. If you look at variables together, like maybe not just signal one, but signal one and 100 together, you get a different picture, i.e. when Scottie and Michael played apart. Our position is you have to compute and use as many different, conjoined probabilities or if you will combinations, as you can, to try and make the best possible list to borrowers of what they need to improve to get a loan in the future.

Peter: Okay, so something that you would do is say, right, you don't just improve these one/two things, you maybe improve these six or eight things because they're in combination. I guess maybe we can...maybe that is part of what you're offering and I want to get into this in a little bit, your Zest Automated Machine Learning product, but is it part of the output or part of the offering to be able to create these adverse action letters that make more sense or that make sense based on what the borrower wants to actually hear?

Douglas: Yeah, that's a great question. Maybe just for clarity, I'll take a step back and say what do the demo tools actually offer and once I get you that maybe it'll be a little bit easier to talk about...

Peter: Yeah, sure.



Douglas: So there are basically three kinds of steps in modeling. There is the actual math itself, there is convincing yourself and those around you that the math is actually correct and there is putting it into production where it lives for however long you are comfortable with it living. You can get from Zest, if you want, a modeling environment in which you build models, but you can also use your own modeling environment or any number of other tools out there.

We do sell some AI little probes, little kinds of measuring things that pay attention to your model build and we use that AI to help you do something which is super hard which is convincing regulators of what you're doing and why you're doing it. After that process we build a model for you, an AI model to look at your AI model, they help you understand and build the model that for any individual applicant will yield the adverse action letters. You've got this AI running to look at AI and the AI and the AI that is looking, is what you ask for help when you deny an applicant. And then we built a bunch of technology when you put your model into production to use other kinds of AI to notice if your model is consistent and [inaudible].

One of the things which is true about many large machine learning models in production is that they're a little bit kitschy, their answers switch a little bit and sometimes kind of for no reason they go running for a cliff and throw themselves off (Peter laughs) and people who haven't read a lot on machine learning models might sometimes forget that and your model runs off a cliff and crashes to the rocks below and you're still counting on it giving you loans.

Our clients get a set of tools that maybe can't prevent the model from trying to throw itself off a cliff, but they tell you, hey, something bad is happening, stop making loans before the model goes south. And also allows you to track things like monthly economic impact and return, lots of sort of useful stuff to the business folks so ultimately, what we're trying to do is build ML-based tools that are invisible generally, they're sort of half in the background, but yield all this stuff that business and audit and regulators can use to run the business and not care anymore that the technology underneath has changed.

Peter: Right, right, got it. Okay, so let's talk a bit about business cases. It would be great to hear maybe a couple of examples of...there was a big name that came out, I think it was just a few days ago, with you guys in the Wall Street Journal, but just maybe take us through a couple of examples of implementations that you've done with names we'd recognize.

Douglas: Sure, so I think the Wall Street Journal, a while ago, outed Synchrony as a client of ours, so that's an old one from like 18 months ago, 24 months ago and in that they described a very large amount on the order of several tens of millions of savings so that's like a name you'd recognize and it's the same process, right, the model was really good, their models were amazingly good and then this kind of monitoring stuff that we did.

The announcement you were referring to a few days ago was with Discover, obviously again, an amazingly talented organization. We're super proud to be their partner and they're using our tools to basically, their goal was to get a handle on increasing losses and our tools have helped them get a handle on that and also helped them interact with the regulators on that point and



they're in production or moving to production to do the same thing, around like is my model okay, is it healthy, etc.

So you get this case where they're obviously quite talented modelers. The business case they were focused on was, hang on, we're seeing some...a little bit erratic increase in losses, certainly not disastrous, but not the trend you like. So they built models on our tools, we helped them build the regulatory and audit documents for those models. When I say we, I say like this AI, which was monitoring their model build, built these things for audit and regulatory and then the model is in production with our other AI kind of bounding it to keep it away from the cliff and the rocks below.

Peter: Okay, so then when you're going into these companies, is it...someone like Discover, do they already have like AI models in place, or they're just sort of using...I imagine there's still plenty of financial institutions using sort of more of a linear regression type models even today, I would suspect, so what is the process like when you go in and I presume you can work with someone who's got a very advanced model and somebody who's got a very rudimentary model, is that correct?

Douglas: That's true, and lots of financial institutions still use a linear regression or logistic regression, some actually just use a score from a credit bureau so that's sort of like the minimalist approach. We find that across the spectrum, from there all the way up to financial institutions that are actually working on machine learning themselves that the tools that we provide add a lot of value; we're just describing them in a different way. So if you're a financial institution that's just using a score, you don't have a lot of data lying around because you're just using a score and so we might not be able to build a model or you might not be able to build a model day 1 because, you don't have the data.

You might need to do what's called a retro and get data from a credit bureau, or you might need to do a testing plan to generate some future loans and allow a model to be built so that might be a different process. If you are using a logistic regression you've got data lying around and some companies that use logistic regression for underwriting also have other interesting data like...maybe they have phone call data, maybe they have collections data that isn't currently being used in their underwriting because it's very hard to use that data in underwriting. But, machine learning models and the models that people build with our tools can use that kind of data.

So for those folks who might say, yup, give us your underwriting data and oh by the way, let's get some of this data from underwriting, excuse me, from collections, let's get some over here from call data and let's use all of that stuff to build a model, which gives us much more information, much more signals; the models are better.

Finally, the organizations that are building their models today, often have really, really great sets of data and sometimes they've got data lying around that isn't used, etc. and where our tools add value. It's helping them get that out of the research lab through the regulatory folks and in



the production. So if you think about it we're sort of adding three different kinds of value which map roughly to the three parts of our software.

Peter: Okay, so then just on that, someone comes to you and they're just using a FICO score and that's pretty much it...so you don't have like an off the shelf model that you can provide them or do you? You said they need to go out and sort of get data from credit bureaus, whatever, so you can build the model. So you don't have something off the shelf that you can just provide them?

Douglas: No, we don't provide generic scores, it's one of the powers of places like FICO that they have these really pretty generic scores. We don't have generic scores or generic models, what we have are the tools to help financial institutions move into the machine learning future. Where presumably they're going to use scores from great places, but they're also going to have more complicated models around those scores.

Peter: So then are you finding more and more institutions already have decent machine learning models in place that you can certainly improve upon, but they're already doing things reasonably well, or in the financial space are people still pretty backward for the most part?

Douglas: I think almost all banks have at least one person or one group doing research on machine learning, but they're mostly focused on the math today so getting from that research to production is quite hard. That said, a huge number of financial institutions have, high quality fraud institutions notably like FICO's Falcon which are machine learning products, they just don't look like it to the banks. So in the underwriting space there has been not a lot of progress in figuring out how to get the mathematics from the laboratory all the way into production.

Peter: Right, right, that makes sense. So you see a lot of partnerships now that are going on between banks and fintech platforms, whether you've got Avant or Amount now doing...they've announced a couple of decent sized banks that they're working with...there's obviously Upstart and others that are coming in as well. How do you sort of play into this shift that we're seeing between fintechs and banks, the partnership model that we're seeing?

Douglas: I think it's important to differentiate kind of a white label kind of approach with a tools based approach so you look at what some of the great partnerships have been. Both Avant and Upstart are really great companies, Upstart founded by Dave Girouard, who I was privileged to work with at Google for many years. So those guys tend to focus more on a white label approach that says, hey, come over here, write me a check and I will put up an underwriting engine for you and that serves certain kinds of needs extremely well, that's not what we do.

What we do is give you the tools to build your own models so in some sense, if you want to think about it in, you know, aphorism terms, a white label partnership is I'm giving you a fish. What we do is we're teaching you how to fish. Each one has really big advantages and really big disadvantages. We've proposed that we want to create organizations, develop an organization, grow from a handle on the side of a financial institution and the easiest way to do that is to actually provide tools. Now we have to provide training and support and a bunch of



other stuff as well which is the downside, but that's why we've chosen the way we've done deals so our partnerships don't look like the kind of white label deals that you've seen a lot in the press.

Peter: Right, right, got it. So we're almost out of time, but a couple of things I want to hit on before I let you go. I'm curious about how you interface with the regulators. You know, there was certainly...a lot of attention was paid when Upstart got their no-action letter from the CFPB because that sort of was the first movement by a regulator saying that, hey, this is...we're interested in actually, you know, exploring and accepting some of these new changes, but some of the things that you're talking about...I'm curious, it seems like it's beyond some of the things that the no-action letter talked about so tell us about your relationship with regulators, how you're engaging with them.

Douglas: So we spend a lot of time with the regulators because ultimately, our clients need assurance that the regulators understand and I've done things like...I've done machine learning training for all of the members of the CFPB, I did ML training for the examiners of the OCC, same in the Fed, same in the FDIC so I spend a lot of time with the regulators. In fact, we drafted a set of, you know, addenda to the FAQ used by the examiners of all three of the major bureaus to sort of make it a little bit more possible for them to ask the right questions about machine learning and get the right answers.

The short answer is everyone believes that ML is the future, and everyone we've talked to believe that ML could be the present, but they don't view themselves as the (inaudible). There are some things that we have to work through. My perspective is that people who say that the regulators are inherently blocky and they're really retrograde actually aren't spending any time talking to them. There are amazing folks who are working at the regulatory agencies because they believe serving the US government is an honor and that protecting consumers is an important moral duty. These folks are not trying to take us back to the 70's, they're hoping to prevent 2008 again. So my experience has been actually quite positive and quite frequent, I'm in DC a lot.

Peter: Right, right. Okay, so last question then, I'm really curious about where you think this is going, I mean, obviously, things have changed a lot in the last five years and I think in the next five years, I imagine, it's going to change even faster so take us through where...how good can these models get? Where will Zest be and sort of your abilities be in five years time?

Douglas: Five years to a startup is an infinite window (Peter laughs) so it's quite hard to predict five years in advance. I think it is hugely indicative that Discover went out with their press release saying, hey we're partnering with Zest. I'd say that the classic big monster dominant quality institutions are starting to pay attention and I expect more of those and for various reasons I expect there will be more of those.

So I think the next thing that happens is how does mass adoption happen, mass adoption for scores took about 20 years. I don't think it's going to take that long for machine learning



LEND ACADEMY

because the process of actually adding the systems and the skills is easier now than it was then so some amount farther. My hope is just that in five years I have a pool table at my office that works (Peter laughs) so that's my big goal. (Peter laughs)

Peter: Okay, well on that note, we'll leave it there, Douglas. I really appreciate you coming on the show today.

Douglas: Hey, thanks so much for having me.

Peter: Okay, see you.

You know, I think most people in credit today recognize the importance of AI tools and how it can really help increase loan performance, it can help expand the credit box and, you know, can bring credit to more creditworthy individuals. It's a testimony, I think, to the work that Douglas and his team have done at Zest Finance that they were able to bring in such large financial institutions such as Discover.

A quick plug, if I may, we have the CEO of Discover and Douglas on stage at LendIt Fintech USA in a couple of weeks. If you're listening to this after April 9th, you'll be able to watch the video on our website, it's going to be one of our featured keynotes of the entire event.

Regardless, I think, more and more of these large financial institutions are going to be deploying these sophisticated AI tools and I think that will be a great thing as we continue to expand access to credit to the underserved.

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

Today's episode was sponsored by LendIt Fintech USA 2019, the world's leading event in financial services innovation. It's happening April 8th through 9th, at Moscone West in San Francisco. It's going to be the largest fintech event held in the Bay Area in 2019. We'll be covering online lending, blockchain, digital banking and much more. You can find out all about it and register at lendit.com.

(closing music)