

AI and You

Transcript

Guest: Dan Turchin, part 1

Episode 122

First Aired: Monday, October 17, 2022

Welcome to episode 122! A few weeks ago, I was a guest on the popular podcast AI and the Future of Work, which really goes into the business of AI talking with people who make AI and IT work in real businesses, real money at stake; and the host and I really hit it off. We're doing the same kind of thing for the same kind of reason and we said, we ought to talk some more, this is really interesting, so the first thing I did was invite him to AI and You, and here we go. My guest today is Dan Turchin, an entrepreneurial leader who is passionate about building great teams that build great products that solve hard problems that change lives. He's doing that currently as the CEO of PeopleReign, which automates the lifecycle of service requests, and we'll learn how that works during the show. He's a serial investor and entrepreneur, having run AI-centric businesses like InsightFinder and Aeroprise, and been a senior director at BMC and ServiceNow.

I've said many times before that the people I want on the show are the ones with a passion for what AI can do for people, and you're going to hear that here and you're going to hear why it makes a difference. Let's get into the interview with Dan Turchin.

Dan, welcome to the show.

Peter, I'm looking forward to this conversation. Thanks for having me.

So, you have had a history here of working with AI in different forms of looks like customer service, perhaps you can take us through how you got into that and what it means to you, what the motivations were, because I'm always fascinated by how people started out in this field and what makes them go.

Yeah, Peter, roll back the clock 25 years. I first entered the space, what I call applying AI machine learning to the field of employee service. Now, 25 years ago, seven companies ago, we did a lot differently. We couldn't even dream big enough to think about the kinds of technology and the kinds of problems we can solve today. Today, I'm the CEO of a company called PeopleReign, SaaS platform that automates HR and IT employee service for large organizations, think McDonald's or the federal government or Walmart or Boeing: Large organizations that are committed to delivering exceptional employee experiences. And ironically, they're looking to use automation technology to humanize the experience that their employees have dealing with internal service providers, because oftentimes, employees would rather interact with a smart virtual agent that learns continuously and speaks their language, rather than being treated like a ticket. So, we're actually using some of these modern technologies. We built a neural net around about a billion historical IT tickets and HR cases. So, its ability to detect intent and resolve employee issues, vastly outperforms the ability of a traditional live agent in a call center. So, we

really think about this as not just the future of work, but what it means to augment the capabilities of humans with the benefits of AI and machine learning.

Now, you said internal service providers? So, is this talking to customers of those companies? Or is it those workers inside those companies dealing with some issue that they have internal to the business.

I can see employees of those organizations dealing with everything from hardware and software issues with their personal equipment to identity access management issues, password resets, PTO questions, benefits, payroll, things that in the ordinary course of business, as an employee, create friction between you and your employer, we like to say, we give every employee back an hour a week, to do more of what they love doing and less of what they hate doing, which is sitting in a queue and being treated like a ticket. So, that that kind of productivity dividend multiplied by a billion employees around the globe, translates into what we really feel like puts a little dent in the universe in terms of using AI and machine learning to ultimately make us better versions of ourselves.

And so, I have a lot of experience as someone inside a large organization, as many of our listeners will, of dealing internal service providers, people providing IT doing exactly the same sort of thing that you're talking about password resets, that sort of thing. And so I'm familiar with things like remedy tickets, ServiceNow tickets, incident reports and ticket numbers. So, make it real for me how my experience would differ with a system like yours? What's the essential nugget of what you do that makes that different?

So, historically, as an employee, when you have a problem with your technology, you have a question about your benefits, etc. You do one of two things, you either send an email to the service desk or you call the Service Desk on the phone, and you rue the day when you have to do that. I mean, you just you try everything possible to avoid having to call the service desk because it starts kind of an infinite loop that even for the most common request, it can span days and it's so dehumanizing to feel like you've just been reduced to a number, and you're waiting for that update delivered via email. That's the old world of work. When we talk about the future of work, we'll always talk about using AI and machine learning to augment the capabilities of humans. So, when you have a smarter service experience, whether it's fully automated, or the live agent who you talk to on the phone has, a set of next best actions delivered to them, as soon as your call was received, that helps them resolve your issue the right way, the first time, man, you feel great. You like your employer, you know that. That day of kind of feeling demotivated and disengaged, gets turned into a day when you love where you're working, you're doing more of the things that you were hired to do, and less than the things that you hate doing. So, you've lived it, you mentioned you know whether it's a remedy ticket or a ServiceNow ticket, what people are in is doing is turning those kind of traditional dumb service experiences into augmented intelligence service experiences.

I'm biting my tongue to avoid getting in trouble in some places. So let's say I've got one of those problems, you can pick one but I engage with the helpdesk am I in this case, talking over the phone to a bot, or am I texting a bot, which one provides the better experience?

Yeah, so your first port of call will be whatever channel you prefer. Ordinarily, whether that's a voice channel, a synchronous chat channel, an asynchronous channel, like email, only now, instead of engaging with a live agent, first waiting for and then engaging with a live agent, you're now first engaging with a virtual agent, and it's using natural language processing to have a more fluid, conversationally relevant discussion with you than what you would traditionally have with a live agent. I know that's a positive, because I know that's a little bit ironic. But what employees using a modern virtual agent, like the people in virtual agent realize is that they would actually rather engage a virtual agent, it's optimized for resolving employee service issues, then the old way of talking to a live agent, it's not really an expert, and has to take your valuable time to learn how to fix your problem. They're not all-knowing beings, like we'd expect them to be. They're humans like us and they need some time to process and research. While the virtual agent, it's always available, it speaks 27 languages, it's trained to respond to 5 million common work concepts. So it's much better than the traditional experience.

So, the difference is that the agent has absorbed a larger knowledge base than a human can. Is that it?

Yeah, exactly. It's kind of like the amalgamation of everything that every live service agent on the planet has ever learned compressed into one digital brain.

And so, it's got to know the knowledge base for that company, right?

Yeah, so it's a combination of what we call an upper ontology, that's kind of a structured vocabulary built on customer-specific content, policies, knowledge articles, appended to what we call a lower ontology, which is generic concepts that span every enterprise, whether that's tax policies or software configuration issues. When you build the upper ontology and fuse it onto the lower ontology, you get the most intelligent service experience possible.

Right. So, why restrict this to internal, why not go for the customer service?

Something I learned in the '90s, when I first got into the employee service space is that every organization is looking to deliver a consistently high quality of IT and HR employee service experience and so if we are to build technology optimized for the "best way to deliver IT and HR employee service," we can do it once and we can transfer that learning across every organization on the planet. If we're instead supporting external customers, specifically, in a data-driven way, using AI and machine learning, we no longer have that data advantage. We need to build separate models for you know, let's say it's Boeing responding to Southwest Airlines about their version of a 737. That's a very different technology than let's say we're supporting Levi's customers supporting jeans returns online through Amazon There's very little overlap in terms of those two technologies. But if we're supporting the Levi's employees and the Boeing employees, we can pretty much build a universal AI engine, that's as relevant for the Levi's employee as it is for the Boeing employee.

I'm not sure I understand yet. Looking at these examples, the core of what you've been talking about would be, for instance, the Boeing employees calling up the Boeing IT desk and saying "I can't get on the VPN" and next to that you'd have B2B, like Southwest, calls Boeing and says, "What's the service interval on this engine" or something like that. And that would seem to be in the same arena, because it's a much more restricted domain of application, it should be easier to understand what they're talking about, they should be, let's say more expert kind of customers more willing to deal with a machine and more capable of having an interaction with it that would be productive. Compared to someone calling Southwest Airlines and ranting that their baggage didn't make it. So, it seems like a continuum, right? Is it a continuum?

Well, so as the Boeing employee calling up asking how to get my Cisco VPN client configured from a MacOS, I'm going to have the same question, whether I'm at Boeing or at Walmart or at Levi's. If I'm an external customer, calling the Boeing customer service desk, and I'm asking how do you replace the air filter in this, 737, that's very different from a Levi's customer, calling the Levi's customer service desk, and asking how to return a pair of jeans.

I see. So, it sounds like you're talking then about the knowledge base of the product. You've built in a lot of specialized tasks, like you've actually dealt with these particular ones that are common to internal businesses, like using the VPN, configuring email clients and things like that and you've got all those in and so if you were to take it into some other larger space, where there are questions that aren't in that database, we got a business that's like agriculture, say, and their employees are calling up and wanting to know things like, "Is this fertilizer good in this area for the soil type?" It's not in your database yet, but it's a reasonable question within that business. Can you train it? Or do you say, "Sorry, that's not what we do"?

It can be but we've specifically chosen as a business strategy to focus on IT and HR, because the Monsanto employee is going to have a relatively similar question about how to get their Cisco VPN configured for MacOS as the Boeing employee, or the Levi's employee. But obviously, when it comes to an external customer of Monsanto versus Boeing versus Levi's, those questions are very different, but the IT and HR questions you're going to tend to be relatively common when it comes to just about any organization on the planet.

So, it sounds like your crown jewels, are this knowledge base that you've curated into the way you've tuned it for your AI, is that right?

Yeah, and to expand it into a broader theme on AI in the future of work, my podcast, which you were a great guest on, we talked about AGI - Artificial General Intelligence - versus AI and ANI, Artificial Narrow Intelligence and at PeopleReign, we've explicitly taken the approach of, first of all ,supervised machine learning. So, we train it on historical data. And second of all, I firmly believe that before we achieve anything close to AGI, AI will perform very well in narrow domains and so we're kind of the best example of how AI can achieve the benefits we'd all expect but only in narrow domains where we can make it exceptionally accurate and deliver a great experience, but only in a narrow domain like IT and HR employee service.

Right, and some of the things that are being done in narrow artificial intelligence at the moment, sort of defy even recent expectations about what “narrow” would be and you really have to pick it apart to see how it’s *not* generalized, like the large language model and the conversations that you have with it, we’re starting to get to the point where it is hard to explain why it’s not AGI; and are you actually getting into anything with the large language models?

Only insofar as we optimize the large language models for our domain. So even if you take GPT-3, or the image equivalent, DALL-E, what you find is that they make for great parlor tricks. So, if I want a monkey captaining a boat, in the style of Monet it’s a neat parlor trick and DALL-E or DALL-E-2 is exceptionally good at performing parlor tricks. But we talk about ANI, and we’ve done these experiments, we actually need to feed GPT-3 specific additional data to be able to get it to perform a task like delivering accurate IT and HR employee service that goes beyond a parlor trick and is better than what you could get from a Google search, that actually requires a lot of tuning of a model of an LLM, like GPT-3 if you’re focusing on a narrow domain.

And that’s very interesting, because that provides a window into where the boundary is between the large language models and AGI, because it’s actually, in circles that you’re quite familiar with and our listeners are quite familiar with, getting hard to find that boundary to delineate it in some places and the Turing test obviously a sort of trite example and that all because I think we’re now we’ve sort of skirted around the boundary of that and said that, yeah, that’s useful in some domains, but not in others and we really need better tests that we don’t actually have for telling how fit for purpose, say a conversational bot is in some areas, like the conversational bots that they had on what was it Ashley Madison, the dating site. Didn’t have to be that sophisticated, right? Just said, “Like the look of your picture. Come here often?” and you’re off to the races. So, in different domains, there’s different standards. Obviously, if you’re Boeing, and you want to have conversations about how parts fit in engines, you can’t afford to make a mistake, you’ve got to understand that properly. So, if you project from where you are now, and the evident passion, you’ve got for this out several years, you choose, where do you see this going? And does it cross the boundary into AGI for you?

We’re explicitly not trying to pass the Turing test; I think and I feel like you might agree with me here. The Turing test is a poor proxy for AGI. AGI, the ability to replicate any human task using machines is a well, we can debate the ethics of it but perhaps it’s a worthy ambition, I think we’re far away, I’m in the camp of probably 30 to 50 years away from achieving AGI in that sense of the word. For PeopleReign, our objective is to improve work life for the next billion employees. And so if we constrain our vision, our set of objectives to that, then we’re not thinking about AGI. We’re thinking about what it means to fully automate the employee experience with internal service providers in such a way that the employee truly prefers, interacting with a fully digital version of a service agent and we think we can get there in 24 months, but we can achieve AGI and nor we want to we would never do anything like, I’d call probably what Google duplex did. We don’t want to fool humans into thinking that they’re interacting with other humans when they’re not. We just want the experience dealing with internal service providers to be so vastly superior when you use autonomous technology that you

never want to go back. Because that hour a week that you get back, let's you know, spend more time with your family be a better, you know, parent relative pursues a hobby and so when we're giving back the planet a billion hours a week, to be better versions of humans, we feel like that's a great way to exercise the real power of artificial intelligence.

I agree. Although - and this is a tangent - my experience of the psychology of high-tech companies in particular, is that if you've got people working 70 hours a week for one of those, and you find a way to give them back 10, they'll use it to do more work. But that's another conversation.

That's the end of the first half of the interview, we're splitting it up to make these episodes more digestible. This really illustrated for me how conversational AI is being monetized and is a viable business, as part of the service function of an internal business unit. And here's Dan really fleshing out for us how that is working.

In today's news ripped from the headlines about AI, Bruce Willis has *not* sold the rights to generate deep fakes based on his face. That claim was reported in the *Daily Mail* and then reproduced by the *Telegraph*, who should have known better, and other outlets. But in an article on the BBC, Willis' agent said that there was no agreement between Willis and the Deepcake company as claimed. Willis did grant the Russian company Deepcake the right to make a commercial for the telecommunications company Megafon using his face on a Russian actor. It's a pretty standard sort of deepfake, quite well done, not perfect. But that's a far cry from Willis acting in future movies via deepfakes.

Mostly I bring this up because I have predicted that we're not far off from having the technology to put a digital human actor in a movie so you could have, say, Bogart and Monroe together. It would still be expensive, right now, but you can imagine how that might get a lot of box office attention. Probably the bigger stumbling block would be securing rights to the actors' images from their estates.

Now in the case of Bruce Willis, this story was made more plausible because he has developed aphasia and frontotemporal dementia, which has made it too difficult for him to perform in recent years. That's a very unfortunate thing to happen to an actor and makes me sad, because I've been watching him since he was in *Moonlighting*, and I liked everything he did and he always seemed like a nice guy. The last movie I saw him in was *RED 2*, and that was loads of fun. So you could imagine that he might sell the rights to deepfake him, it just didn't happen.

Next week, we'll conclude the interview with Dan Turchin, when we'll talk about the broader question of how AI is impacting the future of work, which is what his podcast, with all kinds of fantastic guests - and me - is called: "AI and the Future of Work." Look for it on your favorite podcast platform. Also explore his company, PeopleReign, which apparently he named because he was a Prince fan - Purple Rain, get it? - at peoplereign.io, if your business needs that kind of service. That's next week on *AI and You*.

Until then, remember: no matter how much computers learn how to do, it's how we come together as *humans* that matters.

<http://aiandyou.net>