AI and You

Transcript

Guest: Dan Turchin, part 2

Episode 123

First Aired: Monday, October 24, 2022

Hello, and welcome to episode 123! Today we will conclude the interview with 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. Last week we got into how he uses AI to provide internal business IT support functions, which might sound a bit dry, perhaps, but there was a lot of nuance I found fascinating about how you can make that successful and not another annoying chatbot. So let's get back into the interview.

This really, for me is, is poking at the question of, well, AGI is one of the proxies we use for this, but one of the terms we use for it, but we're having definitional issues with that, and the Turing test has its own definitional issues as well. But our guest next week, Mark Lee, who have actually already talked to is working on robots, with the ultimate game of embodied robot that can hold conversations, and we had a live conversation about what it means to understand things and his take is that you have to be embodied to reach an actual understanding of the real world. Now, do you foresee that the expansion, the development, the improvement of what you're doing here necessarily hinges upon understanding of the real world?

Well, first of all, I look forward to hearing that conversation with Mark, I don't know Mark, I did recently have a gentleman named Gordon Wilson, who's the CEO of a neuromorphic computing company called rein and he talked about their quest to build a digital brain, essentially, built of silicon, but with all the capabilities and even the neural architecture of a human brain, and it was very interesting conversation, and encourage any of your listeners to go listen to that episode on an AI in the future of work, but he was talking about early use cases are primarily things like prosthetic limbs. So, with a human brain, you could control your prosthetic limb, because it has all of the neurons and the sensors to be able to essentially integrate with the human with the human system and so in that regard, I think we should be careful when we talk about what it is that we that we want. By augmenting humans, I think that use cases like that, where it's clearly complementing what a human can do, with this notion of neuromorphic computing, I think that's an easier maybe interim goal to achieve than maybe embodying a robot with the ability to have full human life conversations. I applaud mark and anyone else who's that far along in terms of NLP and achieving human level conversations across a wide range of conversations, but at least based on our research with NLP and what I've learned from guests, and that sort of thing, I think it's a worthy goal and we've seen things like - Woebot is a technology company that has about a $500 million valuation, to be able to automate conversations and resolve patients with dealing

with mental health issues. So, in narrow domains, I think maybe NLP is ready. But I'd love to be convinced otherwise, you know, but a robot and body with the full capabilities of having human level conversations feels like a very ambitious goal based on the current state of NLP.

And he is very ambitious in that respect. If I think about my experience with chatbots, and I don't think I've interacted with yours - I haven't interacted with yours, I'm not sure if I've interacted with anything on that that scale of that type. They've been frustrating because I am the kind of person that if I'm calling on the phone, it's because I can't get the answer from the website or Google, it doesn't exist there and the chat bots that I interact with are only based on that type of information. So, and this is, of course, customer-facing chatbots. So, they're there to basically triage callers to route as few of them to live humans as possible. But when I am talking to a chatbot, I would have some concern, of hitting limits if it doesn't fully understand me. I'm one going to wonder, am I going to trigger something where it it hasn't understood me properly? and it's not sophisticated enough for me to hold a dialogue with it to say, "this is what I meant by this, can you say that back in your own words, so that I know you've understood me properly," which is something that I do with life service agents all the time, because of so many places that seem to hire frontline people that have communication dysfunction. And so for me, that would be a boundary that would be a question of can that that chatbot introspect, if you will, enough to answer those kinds of questions so that I am assured that there hasn't been some accident of misunderstanding, of ambiguity. And perhaps you can translate that into your experience of and tell me where that lies in the domain of problems you solve.

One of the key principles of NLP, at least in its current state, is the ability to detect known unknowns. So, it turns out that the smartest virtual agents today, the ones that are most effective at actually resolving employee issues, are the ones that are able to identify what they don't know and create a seamless fallback experience. When they're asked to answer a question, they don't know, the problem that becomes, meme worthy, or gets screenshotted or mocked publicly, is when virtual agents claim to know an answer they don't know. I think that in our personal lives, we've gotten more comfortable, whether it's an airline, for example, or our credit card company, they're getting a lot better at forcing us into defined column happy paths. So, if I want to check the status of my flight, or check the balance on my credit card, or ask, "when is my next payment due?" At this point, we can pretty much rely on an automated system to deliver discrete answers to common questions. But as soon as we get into things that are uncommon, can I refinance my mortgage, and how many points are going to be applied to my loan things that require some intuition or judgment or are less common. I think that's where this, we reach a wall in terms of the current limits of NLP technology. And I think it's incumbent on the organization, your credit card, your airline, etc., combined with AI first NLP first vendors, like PeopleReign to create, craft an experience for the end user that takes advantage of the strengths of the technology without exposing the end user to some limitations.

So, one of the common problems with supervised learning is that it only is trained to produce results in its output set. So, if you trained something to recognize dog breeds, it only knows dog

breeds, you can feed it a fish, and it will come up with some random dog breed as the answer. When what you really want it to say is, I" have no idea what this thing is; doesn't look like a dog." And that's harder, and a lot of places don't do that. To what extent do you get into that to be able to identify the level of confidence in what you think the person has said is low enough that you need to say, "I don't know. Can you say that a different way?"

It's a critical component, and today, the best AI solutions in data science terms use confidence thresholds to be able to optimize - I know this is kind of wonky - but precision and recall. So, I know you this Peter, but just quick background. So we want the AI model to deliver answers the highest percentage of time possible, but only when their confidence in that answer exceeds a certain threshold. So, highly precise, but also we want to optimize for giving precise answers as high a percentage of the time as possible. So, maximize or optimize precision while optimizing recall. Now, in the case of PeopleReign, we've taken great pains to make sure that there's instrumentation built into the PeopleReign administration console, so that every automated decision is under the control of an SME, a subject matter expert who's not a data scientist, but they can go in and kind of adjust gauges and levers to actually be able to, without knowing the data science, be able to control the amount of true positives and limit the amount of false positives or false negatives.

Thanks. There're some things that I want to return to later. But what I'd like to do now is just move on to talking about your podcast. There's a lot of people working in deploying AI solutions, and of one kind or another. Very few of them do podcasts. I think I know one. So, I'm intrigued. What's your motivation? What's the goal? What sort of mission? How did the podcast get started? What itch does it scratch for you?

We're over 150 episodes in and the original vision for inviting AI thought leaders onto the program was - The genesis of it was, I asked a friend of mine who's an author, gentleman named Mark Settle who's written about the future of work, and specifically a seven-time CIO, mostly in high tech companies. And I invited him to have a conversation with my team who's at a previous company. I wanted them to learn from someone who's a role model to me, and has done a lot of interesting research about how CIOs buy technology. And it was such a fascinating discussion about the future of work that I half-jokingly said to Mark, we need to record this and share this with the world. Because what, you just shared this masterclass on the future of technology and the future of work, it's a universal dialogue, I don't just want my team to hear it. And 150 episodes later, it's it continues to be a passion project, because I'm so I'm so intrigued by all the wisdom that is out there, people like Peter Scott and hundreds of other world experts with unique perspectives and when you put them all together in the form of weekly episodes on one show where the theme is very broad AI, and the future of work, you create a library of conversations and insights that I think when you step back after a few 100 episodes and think about you know, the themes. And I feel like it's, it's become a passion project for me and the fact that, a community over 5000 people a week enjoy hearing these discussions makes me so proud and so motivated to keep sharing them, because what started as just my personal curiosity and, thirst for knowledge, turns out a lot of people share a similar passion and that is and always what

will drive the podcast and we plan to continue to publish these great conversations as long be as people are interested in listening to them.

And took the words out of my mouth to this is why I have people like you on this show, because you evidence the passion, it's not a job to you, it's a calling, and that really excites me and to do that in the context of uncovering this enormous thing. Every time I peel a layer off AI it shows another ten. I'll never run out of ways of exploring this that are useful to me and to the people who listen to this and to humanity. And I think that -not to get too bombastic, but understanding AI is one of the greatest challenges if not risks of not doing it, that the human race will ever face and. And so the fact that I can make some difference in that, and that you can make some difference in in doing that is one of the most important things that we can we can do and that confers on me an obligation to keep doing it. So, in the process of doing the podcast, what have you learned as a result? How have you grown? How has it changed you?

I started the podcast because I had deep philosophical questions about what it means to be human and how is our complicated relationship with technology both making us better humans, and hopefully making us ask deep questions about ourselves. And what I'm learning as I get educated from and get to know more experts, is that we're all wrestling with similar questions. I firmly believe that as a species we'll change more in the next 10 years, than we have in the previous 10,000. And I attribute a lot of that potential for change and this cataclysmic shift that's underway, to us having to reconcile what does it mean to be human, when our capabilities, our fundamental capabilities, the things that were never controversial about what makes us human, are being questioned. I say frequently on the podcast, "Anything that can be predicted is better left to machines. But anything that requires empathy, or rational thinking, is better left to humans." It's a big statement and for me, the kind of existential purpose that the podcast serves is, every week, it gives me a chance to hear versions of that question answered by various experts. It's really, it's helped me evolve my thinking and I continue to be very positive about the role that will, that technology is and will play in our lives. But I think it's a complicated topic, and one that at least for me, I need to think about, at least you know, once a week for an hour, I think that's the bare minimum to really be able to have some kind of comfort level with how I reconcile this juncture that we're at as a species.

So, what are some of the ways in which we have hit the boundaries of that, the places that make you think, "well, this is, this could be disturbing to people, this is forcing those existential questions on them," some of the ways perhaps they feel threatened, or that these questions are become unavoidable?

One of the statistics I cite is that according to the World Economic Forum, AI will create a net new 58 million jobs in the next five years and to me, it's very easy to always take the over in terms of being optimistic about the future of AI and technology and humanity, when I think about not job displacement, and not, you know, the bot apocalypse or the negative impact of AI on jobs. Instead, I think about the new world that we're creating with AI. And so, when I you know, try to reconcile what can sometimes seem like conflicting views or answer hard questions about what the future of humanity looks like, I go back to things like that and think about the

opportunity that every human has to be improved by AI, reinvent themselves, reconsider foundationally who we are. And I've kind of gone on this crusade to help people understand at a very pragmatic level, what are some of the new careers that are being invented? What are the new skills that we should think about investing in? And what we're going to do when we get back that one, two, three, half a day a week, to be better humans? I think when we reframe the conversation in those terms, it's very easy not just to be optimistic, but to feel compelled to go out and help others feel enthusiastic and get educated about what what's ahead. So, I think even on the on the dark days, when we're wrestling with ethical issues and things like that, I always feel like there's you know a very grounding set of principles that relate to how AI and technology are making us better humans.

Right, is that optimistic view a reaction against people having difficulty dealing with the issues today? Certainly the media likes to put out all kinds of apocalyptic memes, that's their job, that's their bread and butter. One of the ones that sticks is the possibility of someone's job being taken by a robot, and there's no shortage of information that might encourage people to fear that possibility. So, is part of your job to help people understand what the real questions are to react against that oversimplification and to craft a path for themselves in this future?

That's absolutely - that's well said. I mean, that is my personal mission. Seven companies in, to this journey. The bad news is I'm kind of ruined for ever doing anything else. But the good news is, it truly is a passion project and you know, I feel like as arrogant as it comes across, I was put on this planet to impact a billion lives and I feel like there's a multiplier effect, and whether it's through the podcast or other means through PeopleReign etc. You know, I firmly believe that as the name of the company implies, technology's great, but People Reign. And so, when you do a Google search for artificial intelligence, auto suggests, comes up with some, some pretty peculiar, peculiar, and pretty dystopian themes, about, what people are implying when they say artificial intelligence and I don't know about you, but I'm not going to stop until that gets inverted and the associations with artificial intelligence are positive, because the actual facts, the actual ways artificial intelligence are being used to benefit humanity. So, Peter, our work isn't done. It's just starting. But there's, like you said, in the popular press, some common notions that we have to work hard to dispel.

Exactly, and I think one of the things I pursue is that the technology evolution is more or less guaranteed. A lot of people that I talked to take a position of, "we should limit this" and my position is, "you can't stop it, that's not going to work." But the technology evolution is guaranteed; the human evolution is what I'm working on, that we need to evolve to keep up with this and one of the things that's happening at the moment is AI is triggering so many questions of ethics right now, bias, privacy, transparency, explainability, and so forth. A lot of them are the same problems that companies have had all along, but they're magnified because AI lets them make the same mistakes only faster and at scale. How do you explore those issues, the multi dimensions of ethics on the podcast and in your mission to look at AI and the future of work?

About a third of the guests in the podcasts are entrepreneurs who are introducing innovative AI technologies and one of the questions I always ask is, what could go wrong? As entrepreneurs, we tend to be irrationally optimistic and we tend to over index on what could go right. But I think from an ethical perspective, it's important to really be aware of the data that we're using to train the models to make the automated decisions, and unpack the potential implications of poor decisions, when we're not very thoroughly aware of and vetting not just the algorithms, but the underlying data that's used to make the automated decisions. So, one fascinating conversation I had is with a gentleman named Krishna Guda, who's the CEO of a company called Fiddler, which is an AI explainability platform, which helps other organizations developing AI first technology to introspect the automated decisions and understand the answer to that question, what could go wrong, but I feel like it's really incumbent on us as a community to make sure we're doing what I call exercise responsible AI.

So, Dan, where do you see yourself in 10 years with this?

At the current pace of learning We're growing not personally not related to the podcast, but as a community as an AI community, I firmly believe that in 10 years, the dialogue will be much more mature and the kinds of conversations that, Peter, you and I are having, I feel like will be mainstream. And the idea that AI is doing good and the benefits from AI and the jobs that we're creating and the lives that we're improving far outweigh any kind of potential downside. I think in 10 years, that will be a very comfortable conversation that we'll be having and AI will be ambient in the environment. You know, whether it's through sensor networks, I think there'll be a lot more data - that's a platitude - but we'll also be much more responsible, at how we store and use that data and I think it will be more useful universally understood that we embrace innovation and we embrace in our AI in our lives, because it's helping our aging parents be healthier. It's giving mobility back to the disabled. It's doing things that today when we're thinking about the bot apocalypse, where we turn a blind eye toward, I don't think in 10 years, we'll be able to turn that same blind die because I think the benefits of AI, will be so ubiquitous and the benefits will be so apparent that I don't think we'll be wrestling with the same philosophical kind of questions.

Wow. This is the end of our time, we will have a link to PeopleReign in the notes the podcast to listen to for Dan is AI in the future of the work than anything else, you want to tell our listeners about how to find you and follow you.

You know, it's less about listening to the podcast, which I certainly hope everyone will do. Like Peter said, AI in the future of work is available in your favorite podcast app. Certainly, I hope you enjoy insights from over 150 experts on the topic. But you know, more important: be part of the community participate actively in the ongoing dialogue about how technology is augmenting humans. As a species, I firmly believe we'll evolve more in the next three decades than we have in say, the previous three millennia. thanks to AI and machine learning. As a community, let's commit to shaping that evolution in a way that's responsible, and also impactful. Let's figure out how to use AI for good to benefit kids, those with disabilities, the underserved and the forgotten. Let's use AI to impact a billion lives. I encourage everyone listening to join us on this journey. Start by contacting me, mailto:dan@peoplereign.io. People and reign like the rain of a king or

DM me on Twitter. I'm at the church and tell me about your passion and I will connect you to others with similar interests. Peter, I so love this conversation and I can't wait to continue it with you soon. Thanks for having me.

That's the end of the interview. Here you had Dan, who's running a company to develop and sell AI services, who's so taken with what AI means for us and how we work and collaborate that he starts a podcast as a way of expressing that passion, that enthusiasm, that quest for knowledge about this brave new technology called AI. I mean, from personal experience I can say that one does not start a podcast to make money. Not unless you're Joe Rogan. So it's hard evidence of a labor of love at work. I'm sure we'll be hearing more from Dan in the future. And in the meantime, take in his podcast by looking for AI and the Future of Work" on your favorite podcast platform

In today's news ripped from the headlines about AI, this came up shortly after my conversation with Paul Newman of autonomous vehicle software maker Oxbotica a couple of weeks ago, when he mentioned the problem of avoiding *stuff* on the road, and also mentioned the company Aurora. Well, Aurora has announced that they are teaching their self-driving trucks specifically to avoid stuff, like pieces of tire or sofas. This is important because a AAA study showed that nearly 37% of all deaths in road débris collisions result from the driver swerving to avoid hitting an object. And their software is focusing on this question of things that are not supposed to be in the road. They're hoping to have fully driverless semi trucks out by the end of 2024.

Just a bit of editorializing prompted by recently rewatching the series *Star Trek Continues*, which for all you *Star Trek* original series fans out there, I can't recommend highly enough. It's a fan production, not that you'd know it from the production values, set in the original universe with different actors of course, but perfect sets, and fantastic writing that takes off from the moment the original series ended. And I noticed that the amazing final episode was cowritten by Robert Sawyer, who was on this show not long ago. So look on YouTube for *Star Trek Continues*. It's all done by fans as a total labor of love. They're not allowed to make money off it, and yet they got cameos from some big name actors who were clearly also relishing the chance to appear in a real Star Trek episode set on the original *Enterprise*. Vic Mignogna, the actor who plays Kirk must have signed some kind of pact with the devil to make him as much like Shatner as he is, because he has so nailed the Shat that you'll forget that it's someone else. And part of the reason they are not allowed to make money on it is that they reuse the original music but rerecord it with a new orchestra.

Anyway, what does that have to do with this show? It's because every time I think about what the best possible outcome of the human race in the future could be, I land on *Star Trek*. I'm not the only one – every other futurist I've heard who tried to answer that question with an example also named *Star Trek.* It's not about the realism of the technology, because we already know that the future won't look like *Star Trek* because it will include AI that's already much smarter and more ubiquitous than the ones they had. But the insight I had is that the reason we all have *Trek* as our go-to utopia is that it shows the best possible outcome for how humans have become fully engaged. One of the depressing possible futures—but not the worst by a long shot—is that humans end up having everything taken care of for them by machines and we just stagnate in front of the TV, like in the movie WALL-E. But what would it look like to go as far as possible in the other direction? Well, you'd have humans who were being fully used for what makes us human, being our most compassionate and challenged, and that just describes *Star Trek*. It's not *just* the aspect of becoming the kind of species that deserves to inherit the universe, it's also

how we would find that role being the most enlivening experience for us. Whether or not that's ultimately realistic, it's a North Star for a lot of us, as an idea of what we could do when our needs are taken care of and so many of us are not stuck in the survival level of Maslow's hierarchy of needs.

If you want a metaphor, it's like the difference between a light bulb and a laser. A 50-watt light incandescent light bulb is not particularly light; won't do much. But a 50W laser can burn holes in metal. The difference is that it's coherent light, it's all going in the same direction, it's all in the same phase, it's all supporting each other. And as a metaphor for humans, we're like the incandescent bulb in that so many of us are going off in different directions and fighting against each other and not cooperating or understanding each other.

And that means that to a large extent, we cancel out each other's output. It's a miracle we get anything done at all. If we learned how to cooperate and move in the same direction, to be in phase, in sync; well then, we would be like that laser. And that's, again the Star Trek future. I hope that one day we will be worthy to be the creators of superintelligences by being more like that focused laser.

Next week, my guest will be Mark Lee, Professor of Computer Science at Aberystwyth University in Wales and author of the new book, "How to Grow a Robot: Developing Human-Friendly, Social AI." 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