EPISODE 1346

[INTRODUCTION]

[00:00:00] KP: The dream of machines with artificial general intelligence is entirely plausible in the future, yet well beyond the reach of today's cutting-edge technology. However, a virtual agent need not win in Alan Turing's imitation game to be useful. Modern technology can deliver on some of the promises of narrow intelligence for accomplishing specific tasks. PeopleReign has created a virtual agent for IT and HR employee services. This agent's goal is not to replace a human agent, but to augment them by handling some of the requests and elegantly handing off to a human in other cases. In this episode I speak with Dan Turchin, CEO of PeopleReign, about their virtual agent and the future of work.

[INTERVIEW]

[00:00:50] KP: Dan, welcome to Software Engineering Daily.

[00:00:53] DT: Kyle, pleasure to be here. Thanks for having me.

[00:00:56] KP: Well, tell us to kick things off, just a quick note about your podcast. We'll revisit more about what you talk about there. But for listeners who know you from that outlet, what is your show and what do you guys cover?

[00:01:07] DT: So my day job is I'm the CEO of PeopleReign, the AI first system of intelligence for IT and HR employee service. But my passion project is my podcast, AI and the Future of Work. We've done about 100 episodes, here we sit in September of 2021, and I'm really proud of some of the amazing people that have come on the show and agreed to hang out with me so I can pick their brain and get their perspectives on the future of work.

We tend to host a pretty healthy distribution of entrepreneurs, CXOs from industry, academics, journalists, futurists, kind of across the gamut, and we have some amazing conversations about things in or around the topic of the kind of societal and cultural implications of the introduction of AI and automation technologies into the workplace.

[00:01:59] KP: Well, that seems to align quite well with what I've learned about PeopleReign. Can you tell me a little bit about what the team's mission is?

[00:02:07] DT: Yeah, PeopleReign, we are making work life better for the next billion employees. So we feel very strongly that technology is amazing, but people do and always will reign. And so what we've done is built an AI first system of intelligence that first and foremost augments the capabilities of humans. We talk about AI as being artificial intelligence. I don't think there's anything artificial about it. I think when AI is done right, it makes humans smarter first.

And so in the context of employee services, if I'm the employee at a large company, what I really want is just to have my employer take away all the friction that's associated with getting my tax forms processed, or getting a new laptop, or setting me up back to the office technology issues, all those things that kind of can lead to demotivated employees are actually really good candidates for autonomous call deflection.

And so what PeopleReign has done is built four applications on top of a unified AI platform that all help automate the life cycle of resolving an employee's IT or HR request. And we've done that by amassing the world's biggest ontology or kind of structured vocabulary having to do with how employees of large enterprises request and receive work. And we've used that to automate the process of developing language models for conversational AI, but also models that can predict who should work a ticket, models that can help a live agent fix more problems faster, and models that can predict what issues are likely to occur in the future for employees before they've happened. So four distinct AI tasks automated through the application of this core set of domain ontologies built for employee service. And we take that to market in the form of a SaaS platform, sell it to large enterprises, organizations like McDonald's or organizations like Walmart, Boeing, Dell, organizations that have a lot of employees, they're geographically distributed, they speak all kinds of languages. And one AI platform can go a long way toward making humans the best versions of themselves.

[00:04:32] KP: So let's imagine I'm an employee at one of the clients you mentioned and I have some problem. Maybe I can't seem to get something to print on the local printer. How do I get started and what's my user experience?

[00:04:44] DT: So your first port of call is going to be the virtual agent. So you're going to go into your preferred collaboration tool, and that could be something like a Microsoft Teams or a Slack, but it could also be SMS. It could be a voice portal. It could even be traditional email. And so using whatever tool you would traditionally use to get support from the help desk, you would instead be interacting with a virtual agent. And because in the case of like you need to get connected to a printer, it turns out you're probably not the first employee who has ever had to get connected to a printer. So in fact, PeopleReign has been trained on millions of examples of how you connect to a printer.

So what PeopleReign's going to do is introspect your interaction for context. So who you are? Where you are right now? Are you traveling? Have you connected to the printer in the past? Do you have open tickets related to printer issues? Is there a printer that you own? And based on that, it's going to know a lot about the right approach to take. When it figures out the right set of next best actions to take, it's going to engage you conversationally. Now, if at any point in time you decide the virtual agent isn't fixing your problem adequately, you can always pull the rip cord and escalate to a live agent and then the live agent will have context and proceed where the virtual agent left off. But the real benefit of that AI first system of intelligence is you're going to get connected to the printer faster than you ever would have previously. So you get a great service experience. And from the perspective of the service provider, your company's IT organization, they just delivered a much better service experience with far lower costs than would previously be required. And all of the human agents, if the virtual agents not sufficient, are going to be made a lot smarter by having their intelligence augmented through the system of intelligence. That's how it works.

[00:06:38] KP: So you describe part of that process in which that virtual agent is going to ask me a series of questions, maybe something like have I connected to it before? Or perhaps they could look that up. I don't know if that part's automated. But it seems like there's a decision tree or something like it embedded there. How does that get instantiated? Where does the content for this tree come from?

[00:06:59] DT: Yeah. So you authenticate yourself to the virtual agent. And through the act of authenticating, PeopleReign knows a lot about the context in which the request is being made. So again, PeopleReign knows who you are, where you are based on your schedule, what your role is? Okay, so an approval is needed or maybe something needs to be ordered, it knows your budget thresholds. But more importantly, it can look at assets you own to know what printer you're supposed to connect to. It can make sure that it's not leading you to connect to, let's say, a resource that you're not allowed to access. So just by the simple act of authenticating, PeopleReign has access to a lot of other systems that can provide context before it even engages you in that interactive dialogue.

[00:07:41] KP: Well, I'm used to a certain escalation procedure. If I have a technical problem there's often like a tier one support that tries to help me with what seems to be kind of a known script. And then maybe I get escalated to tier two. It seems like maybe you're replacing tier one support. Do you look at it that way?

[00:07:58] DT: Yeah, I look at it more with augmenting tier one support in the sense that traditional tier one support, first off, in a large enterprise, most often it's outsourced to a tier one service provider. And a traditional contract that an enterprise has with a tier one service provider is for on the order of $22 to $25 every time the phone rings. And those numbers are from the Help Desk Institute. It's kind of industry metrics.

And the challenge is that that traditional way of delivering it service is A, the most expensive, and B, the worst possible experience for the employee, because typically those are scripted interactions and it's very rare that the issue would be resolved at tier one. Unlike using PeopleReign in a system of intelligence whereby the virtual agent is going to fix most common problems the first time, because it's been trained on billions of examples of the kind of problem that you're having. So it's as if you could amalgamate the sum total of the world's knowledge about how to fix printers. Into one digital brain, you'd get PeopleReign. So not only is it a better experience, but it's also about 90% less in terms of cost per issue than the traditional model of outsourcing your tier one. So tier one may still be involved, but the experience is better for the employee, and the same tier one service provider is typically able to take on about 10X the capacity, the volume of tickets for the

same amount of resources, or potentially use the same amount of resources to service 10X the volume of tickets. So either way, the enterprise wins. Either way, the employee wins.

[00:09:36] KP: The promise of learning some of those steps and providing immediate feedback from a trained agent is really appealing. I know there must be tons of logs a lot of these corporations have assembled with human operators, and maybe some of those do require a certain touch. But for all of the very simple how do I reset my password kinds of things, you might have a wealth of training data. And you've mentioned some things about it. Can you expand on what you're able to train on? Is there a well-known public corpus or is this data you had to build as you go?

[00:10:07] DT: Yes is the answer to both of those parts of the question. So a lot of what the AI models get trained on comes from the public domain. So examples, public wiki content, or documentation for Office 365, or Reddit forums. There's a lot of good content out there about how to solve common problems. But keep in mind that every enterprise has its own policies, its own geographies and languages that it supports, its own acronyms.

And so at PeopleReign, part of the unique differentiation of the technology is that we split these what we call domain ontologies into an upper and lower where the upper is trained just on enterprise specific data and the lower is shared. It's generic content that can be shared across all tenants in a multi-tenant environment. So the idea is that any answers that come from the organization's upper ontology override anything that's generic, but the two work in harmony to make sure that wherever the best answer lives, PeopleReign is going to be able to surface it and provide that via this interactive dialogue from the virtual agent to the employee.

[00:11:16] KP: In the earliest days of AI, there were big dreams of encoding these sort of knowledge graphs and networks and having ontological data brought in by some expert who can structure it. And I don't know if there's been a trend away from that, but there's certainly been a growth in a lot more deep learning techniques where they hope everything can be learned from first principles. What's your onboarding procedure like? What's it take to get someone stood up?

[00:11:42] DT: Yeah, well first we'll take a little history tour or trip back 70 years. It was the 50s when the term AI, artificial intelligence, was invented by a group led by Marvin Minsky at an event at Dartmouth College. And then kind of the whole objective behind artificial intelligence was to make computers perform tasks that were traditionally done by humans. And then along came Alan Turing, and we came up with this awful idea that we could make computers pass the Turing Test, which was essentially make machines that were indistinguishable from humans. And that led to going to AI Summer, followed by a deep and dark AI winter. And AI winter was because we had set such unrealistic expectations for what general AI was going to be capable of doing.

Now it turns out, here we are in the 2020s, and we finally realized that to get back to kind of call the new AI summer that we're currently experiencing, what's required is thinking about AI in narrow terms. So while we're far from being able to pass the Turing Test to really do anything that a human can do, in narrow domains like what PeopleReign does, AI can be very effective. So we're able to achieve you know levels of fluency, and accuracy, and employee experiences that meet or exceed what you could get as an employee interacting with any kind of live agent who is unaugmented with artificial intelligence.

So kind of where we're at and the way we do the training and we apply these models, it's only as accurate and it only yields such a great experience because we've very intentionally chosen to focus on a narrow domain employee service. And once you do that, you realize the technology is actually available and the models can be accurate. The machine learning aspect, they get better all the time because the domain content is so specific that we've created this flywheel that's led to millions of employees around the world using PeopleReign to get better experience than they've ever had before. But that doesn't necessarily mean that we're ready to generalize the technology to any domain. But that's kind of a little bit more about our approach and why what we do for employee service is working so well.

[00:14:08] KP: Makes sense. Do you have any metric, or I guess it could vary use case to use case, but the percentage of various questions, or tickets, or whatever is inbound that you can manage in an automated fashion versus those that need to be passed along to a human agent?

[00:14:24] DT: Yeah. So there are five key – Well, we refer to them as KPIs, or key performance indicators, that are used to measure the effectiveness of a service organization. Not necessarily one that's AI-driven, any service organization. And we do have benchmarks with regard to PeopleReigna and AI. The key one kind of the one that you are alluding to is what we call autonomous call deflection. And in a typical environment, about 65% of the common issues, your password resets, your application provisioning, your system configuration, about 65% of the inbound issues are good candidates for autonomous call deflection. So that's a key performance indicator. The other four that service organizations use to measure the quality of their service delivery are your cost per ticket. I mentioned oftentimes PeopleReign's extracting like 90% of the cost out of traditional service delivery. Another one is obviously call volume. So the call volume is going to go down. But the others that are equally important are customer satisfaction. So in a large enterprise, you're typically going to sample a subset of the employees that have submitted a trouble ticket and ask them how their experience was. Your customers stat scores are going to go way up and your self-service adoption, the percentage of employees that turn first to the self-service capability, that's going to go way up. So across all of these KPIs along with the fifth one being your mean time to resolve an issue, how quickly it gets fixed, all five of those are going to benefit from the system of intelligence. But oftentimes, to your point, the leading one is 65% autonomous call deflection.

[00:15:59] KP: I could envision a world in which you do a deployment and a customer is very concerned with a controlled user experience. Not too many surprises. And your virtual agent would you have sort of a rigid thing, a process for what it handles and what it doesn't. But there's maybe also the opportunity for that agent to learn from experience and some of the conversations it can see that get escalated. I'm curious about where you lie both on what you can release or your interest from an R&D perspective on a topic like that.

[00:16:29] DT: Yeah. So very foundational to the technology that we built at PeopleReign is that the employee can't get penalized if the virtual agent isn't sufficient. So we very intentionally talk about automating the full lifecycle of an employee service request. And what that means is that if at any point in time the virtual agent isn't sufficient, there needs to be a seamless handoff to a live agent, at which point it's now the responsibility of

PeopleReign to make the live agent smarter. So that's, again, very intentionally designed into the architecture of the product. There are four applications, and in fact two of the four are designed to help the live agent. So the whole kind of user journey is all predicated upon what's best for the employee to get back to work as quickly as possible.

So if you think exclusively through that lens, then you design a system that's very different from what else tends to be out there, which is largely chat bots. And we think chat bots are great if you want to see a parlor trick. A virtual agent you know respond appearing to understand the nature of the question. But we think if the metric is the employee experience, optimizing it, then it's really incumbent on the system of intelligence to wrap a layer of intelligence horizontally around the whole process and only “get credit” when the employee says, “Yep, that fixed my problem.” So the whole life cycle involves the employee being in control of their own experience and PeopleReign not kind of planting the victory flag until the employee says whether it's to a virtual agent or a live agent, “My problem's fixed. I can now get back to work.”

[00:18:10] KP: There's a wide range of IT requests that could come in, and often they're placed in by, I guess, my own experience looking at some of these, I was never really in a support role, but I’ve seen a lot of support tickets where you simply could not establish what the problem is from the sub-sufficient amount of details the initial ticket was put in. What does the virtual agent do in that situation?

[00:18:30] DT: A good 35% of the time, the question comes in and it lacks context, or the utterance, the way they ask the question is ambiguous. And in those cases one of the core capabilities of any virtual agent needs to be understanding what it doesn't know. So the known unknowns problem is actually a really hard AI problem to solve.

And the way the PeopleReign technology assesses when it “knows” enough to answer the question versus when it doesn't know enough is by calculating a confidence score with each reply. So there's actually a full API surface area over PeopleReign. And when that utterance is passed via API to PeopleReign, it's going to respond not just with the most likely answer but also with the confidence threshold. And the customer of PeopleReign is going to set a minimum confidence score below which the virtual agent should not respond.

So for example, an employee might come in and ask what's on the menu in the cafeteria? A perfectly reasonable question. Not for an IT virtual agent. But rather than giving verbal abuse to the employee or slapping them on the wrists, the people are in virtual agents going to come back and probably it's not going to have an answer that it's very confident in. It doesn't know about menus and it doesn't know about cafeterias. And so what it's going to say is I'm not sure how to answer that, but can I bring a live agent on the line to help you resolve that? Or it might detect that it's a better fit for, let's say, an employee culture virtual agent, or an HR virtual agent, and hand off to a virtual agent trained specifically to be able to answer that kind of question. So yes, about 35% of the time, the virtual agent won't have a suitable answer. But in a hundred percent of that 35% of the time, it will come back with an answer that's not gibberish and be able to hand off with context and hopefully still resolve the employee's issue.

[00:20:34] KP: We've talked through some of the IT support examples. Could you give me some examples of HR requests that a virtual agent can field?

[00:20:42] DT: Yeah. So top three types of issues that enter into the HR ticketing inbox relate to payroll, PTO and benefits. So oftentimes, whereas in the world of IT, most often the tickets are related to break fix. Something's broken, I need it fixed. In the world of HR, it's more often how do I or FAQ kinds of questions. So typical issues are when will I get my tax forms? Or how do I enroll in the benefits program? What's our gym subsidy? So things that can be answered in typically a one-shot fashion where there's a discrete answer to a discrete question.

From an AI perspective, the thing that's challenging about being able to address those kinds of HR issues is that oftentimes they're organization specific, and oftentimes the employee's use of language is very colloquial or informal. And so there's a greater need for those domain ontologies to be able to kind of disambiguate the utterance and look at not just concepts, like PTO, or benefits, or payroll, but relationships among concepts as well as all of the different ways that those same questions can be asked to be able to deliver the right answer the first time. So different set of challenges, different set of questions, but a well-trained virtual agent can be just as effective in that domain as it can be in IT.

[00:22:15] KP: On the IT side, it seems like bad advice is useless advice. You wouldn't come with a problem and be told how to accidentally reformat your hard drive. You would just be told something that wouldn't work. You'd be frustrated and escalate. On the HR side, bad advice could be wrong information. Are there any extra ways you have to pay attention to the context and even the presence of double negatives that are tricky to make sure you're not giving the employee the wrong information?

[00:22:44] DT: Yeah. So, Kyle, to your point, I'd say in both domains, the value of a wrong answer, a false positive, it's not just negative, but it's worse than useless. It's worse than not having any recommendation at all. In the IT case, it can lead to you know bricking your laptop. Or if you install the wrong antivirus patch, for example, and you expose yourself to malware, like the implications can be pretty severe. In the HR context, there is PII, personally identifiable information, involved. And oftentimes there can be issues related to, let's say, mental health or things that have health impact. And so in both cases I'd say that the impact can be dire if the virtual agent goes awry. And that's all the more reason why these confidence thresholds are so important and we're able to – Using the machine learning principles of precision and recall, we're able to tune the models. And the customer is actually provided with instrumentation to do this on their own. But to make sure that in high stakes situations like examples I gave, mental health recommendations for example, that the models can be tuned so that thresholds are very conservatively applied. So if in the context of a mental health question the model may only – Anything below, let's say, I don't know, an 85% confidence threshold, and the system will recommend that it not provide an answer in domains like that versus let's say when's the open enrollment period. The stakes are lower for that, but that's the whole value of using this kind of precision recall approach that we're optimizing for making the right recommendations as often as possible, but not at the risk of providing false positive answers.

[00:24:34] KP: When it comes to HR questions, maybe this is true of some IT questions, but I can't think of any. But HR, I might have a question about leaving the company that I was uncomfortable asking a superior or a person. The virtual agent is kind of appealing for just understanding what my rights are without raising alarms or things like that. How do employees view privacy around interacting with the virtual agents?

[00:24:57] DT: So privacy is definitely one of the bedrock principles of we call responsible AI. Making sure that the employees are fully aware that they're interacting with a virtual agent. Any decision that the virtual agent makes on their behalf should be transparent and explainable. So no black boxes, and the system needs to be configurable. Like I've been talking about, configurable confidence thresholds, so that there's zero likelihood that the system will go rogue.

Now to your question, if I'm asking a question that I wouldn't want, let's say, a live HR agent to be aware of, that's a situation where we recommend the virtual agent not be used. Because in almost every situation, anything that the virtual agent can supply to the employee also is visible to the live agent. And in fact, unless the system is configured differently, all of the transcripts from the virtual agent discussions are tracked so that the organization can see what questions are being asked. So there are other ways to configure a system to maintain the employee's privacy. But the traditional way to set up the HR-based virtual agent is not to have the employees ask questions that they wouldn't want their employer to know about.

[00:26:15] KP: Do you find that there are any insights or learnings that come out of reviewing some of those transcripts?

[00:26:22] DT: The biggest benefit of a system of intelligence for the organization using it is that they unlock the intelligence trapped in their data. So I'll give you an example. So oftentimes self-service initiatives at large enterprises fail because no one person in a knowledge management, knowledge management being a function that produces the content that's used to answer the questions, no individual has great visibility into all the ways employees ask questions or the fact that the questions change dynamically.

Last week before we had the, whatever it is, iPhone 13 or something, nobody knew to ask questions about the iPhone 13. And now there are a ton of questions about carriers, and settings, and APNs, and various kinds of things that obviously there's no knowledge associated with. The benefit of a system of intelligence is it continuously learns. So like I said earlier, it's kind of the amalgamation of everything that anyone around the world knows

about a particular topic. So the real benefit to the organization is that you have this new source of digital labor that's continuously improving. And by continuously improving, it's making everyone else in the organization better. That in and of itself is really the ROI, or the return on the investment, in a system of intelligence, because without having to invest with a fraction of the investment in retraining and up-scaling a new content, the system of intelligence is basically guiding the process of making sure the best answers are always available. And that's what leads to a better employee experience. That's what leads to better self-service adoption. That's what leads to better customer satisfaction scores, and really an overall productivity dividend for the business.

[00:28:07] KP: If I were a data scientist inside a company and it was my responsibility to help with an integration like this, it would be very appealing to me to know that there's some configuration options you were mentioning so that I could do any fine-tuning to better suit the organization's specific needs. Yet there's a fine line between an elegant dashboard of parameters and something with so many knobs that no one knows how it works. Can you speak at all to the design challenges in coming up with the ideal solution you want to deliver for customers?

[00:28:36] DT: Yeah, our target persona for the PeopleReign administration console is the administrator of the underlying system of record. So the system of record for IT and HR tends to be one of ServiceNow, Workday, Atlassian, Sales Force, Fresh Desk, something like that. And whoever is administering that underlying platform should be capable of also administering PeopleReign the system of intelligence.

So what we've done with that persona in mind is we've provided instrumentation that shows things like if you turn this lever one notch to the right, you're likely to get you know eight percent more questions answered, but potentially an additional two percent of those questions might be answered in a sub-optimal way. So there's a sensitivity analysis based on some really hard math and statistics on the backend, but it's surfaced in the form of these levers and gauges for sysadmins who don't have any data science background. They just understand kind of workflow and ticketing process. And we think that that's one of the big advantages of using PeopleReign is that we explicitly don't expect there to be any kind of data science resource in-house. All we expect is that someone familiar with the business

process, an SME, or a subject matter expert, can extend their work managing the system of record to managing the system of intelligence.

[00:30:08] KP: It would be totally unreasonable for anyone to think your virtual agents should pass a Turing Test, of course. But by picking a subfield, just some niche areas that that's what the virtual agent is going to be good at, you have an opportunity to be narrow general intelligence. Yet I'm curious, from a user experience point of view, do I feel like when I interact with the agents that there's a sort of narrow intelligence, or is it more formulaic for the purpose of just getting workflows completed?

[00:30:39] DT: It feels narrow in the way it feels narrow to call the help desk and talk to someone who's a specialist in IT, or HR, or finance, or legal, or whatever it is. When you call that help desk because you need an NDA approved and it's a legal request, or your laptop broke and it's an IT request, you don't expect to be able to have a conversation about how do I cook a turkey? Or what's the cheapest flight to Salt Lake City.

Similarly, when you interact with the virtual agent that's trained to respond to IT and HR issues, it's in a place in your enterprise service portal where you go when you have an IT or HR issue. So in fact, the experience with PeopleReign is much smarter than what you'd get from a traditional call to the IT or HR help desk, because at least if you ask something that's gibberish from the perspective of the service that it provides, it'll know what it doesn't know, and it'll provide kind of a safe landing place that involves escalation to a live agent. But at no point would it be worse than the experience of having that interaction with a live agent dedicated to that domain.

[00:31:52] KP: There seems to be a breath and a depth of possible avenues you could be exploring in the future. There are other departments you could help or other industries you could get into. And a depth in terms of the amount and variety of questions you could be answering. Where are you trying to push forward?

[00:32:10] DT: So the logical adjacencies are more – We refer to as horizontals as opposed to verticals. So verticals would be specific industries. So AI for healthcare, finance,

manufacturing, higher education, etc. So PeopleReign has very intentionally stayed horizontal. So IT and HR are things that every one of those verticals has the same need for.

So similarly, the adjacent areas beyond IT and HR are things like security. There's a high overlap between IT and security. In fact, oftentimes, employees don't know where to go when they have issues related to access or patching or things like that. Similarly, there are adjacencies to HR within a large enterprise that are also horizontal, things like operational support, or I mentioned legal. Things related to finance often have an overlap with HR. So those are kind of adjacencies that aren't directly related to IT, HR, but oftentimes the employee requests bleed from one of those into one of the others.

[00:33:16] KP: Do you have a vision for what the interaction that employees will have with work in the future? I mean, if we're going to have better and better technology giving us the libraries and tools to power virtual agents, it seems like they'll give us more and more services. What will be my share of human versus agent interactions with an employer in the future?

[00:33:38] DT: The future of work is ambient intelligence. My work space, my workplace, is infused with this layer of intelligence that today gets kind of embedded in a virtual agent. That essentially gets kind of built into the fabric of the enterprise. I'll give you an example. When I am traveling and I walk into a new place of work, it recognizes me. To your earlier question about the printer, it knows what printer I should connect to. It gets me on WiFi. If I need my badge to be updated, it updates the badge but just for the period of time that I'm in the office for. It looks in my calendar, it knows who I'm inviting to my meetings. It gets them on the network. So all the things that tend to generate today, the questions to the virtual agent or to a live agent, all of those things are going to be automated. So before I even have the request, the system of intelligence is there for me. Kind of like this kind of digital concierge that follows me around the workplace.

And from the perspective of the service provider, we see a time when everyone that's in a service delivery function is going to essentially get back an hour a day 12.5% of their lives because automation is going to lead to better answers faster and it self learns. And the challenge for us as a race of species is really thinking through you know what does it mean.

And maybe it's five, maybe it's ten years off, but what does it mean in terms of the future of humanity? To have that extra hour a day, how do we use that to be better people, spouses, siblings? We can pursue hobbies. We can pursue philanthropic initiatives.

From where I sit today, because I see that future barreling at us pretty quickly, I feel like kind of the enduring impact of PeopleReign will be helping make humans better by kind of embracing this partnership with machines and becoming the best versions of humans that we can be. Kyle, when you and I are having a version of this conversation in the next five, ten years, that's going to be the starting point.

[00:35:47] KP: Well, if I were to engage my, I guess, deep paranoia and think about a knowledge worker who's remote and their core contributions are really over slack and emails and maybe some GitHub commits, they might quickly realize all of their contributions are digitalized. And if you have very, very, very optimistic views of what AI could do, you could have a worry that you're training yourself out of a job. What do you see as the reality for someone who has those fears?

[00:36:14] DT: Many have that concern. And I would allay that fear by saying, if today, you're in a profession that is that easily automated by a machine, it's probably a good time to think about up-skilling or re-skilling. I firmly believe that, in the future, anything that can be predicted is better left to machines, but anything that requires judgment or empathy is better left to humans. So now is a great time to be thinking about that partnership that humans will have with machines and thinking about how to embrace a career, how to prepare for a career, where your colleagues will be machines and you'll want them to be your colleagues because they're going to make you better at what you do. But they're never going to replace what you can do as a human when it comes to coaching other team members, exercising empathy, judgment, being a colleague that you'd want to go to to have a heart-to-heart conversation. These are things that even as we look 10, 20, 30 years, 50 years into the future, those innately human skills will never be automated. But it is really important that we think about a world that's not in the too distant future whereby we'll want to embrace the capabilities of machines to make us better versions of ourselves.

[00:37:41] KP: And in that future, do you see companies like PeopleReign being major service providers at the – I don't know if you're first in this or if there are others, but that we might see there are just service providers that offer a supplement to one department or type of profession spring up and we'll have truly virtual assistants for scheduling and maybe one for helping me with taxes. To what degree will my workforce be bots and automated services?

[00:38:10] DT: Yes. So we foresee a future not too far ahead whereby all of the routine tasks that like I said are predicated on predictions, those will be augmented by machine learning models and fronted by things like NLP-based virtual agents. But again, what's going to happen is you're going to see so many new professions, new tasks, new roles for humans introduced.

For example, just in the field of AI ethics, there's going to be a high need for millions and millions of people around the world to monitor how automated decisions are being made and tune those decision making portals, like what I mentioned, the PeopleReign administration console. There's going to be a need to understand data and understand how to explain automated decisions to employees and understand how to take an automated decision and operationalize it. So absolutely, we're entering a world where this fusion of human and machine is going to be prevalent. But I also think we're entering a new world of work where so many new opportunities are going to be available that it's really the ambitious types that today are thinking ahead about the future of their professions and they're the ones that are going to be on the right side of technology, on the right side of innovation, as more automation is introduced.

[00:39:34] KP: We've gone from a time when virtual agents were non-existent to a time when they were kind of laughable or cute, to a time when now they're useful in a lot of situations. It seems like there's an accelerating trend here. What are some of the key technologies that are pushing it forward?

[00:39:51] DT: That's a good way to frame the question in terms of the evolution, and I'd argue that sometimes virtual agent technology is still overused. I think most of us would agree that the experiences that we have even today with our airlines, or our credit cards, or

service products, certainly the cable company, are still pretty poor. I think the reason why we're having so much success in the domain of internal employee service is because we focus narrowly.

So it's partly some amazing technology and use of neural nets and some of the kind of foundational technologies, storage getting cheaper, compute getting cheaper, public cloud resources available, all of those underpinning technology is really important. But I think approaching the problem from the perspective of narrow domains that are good candidates for kind of autonomous or the introduction of autonomous technologies, I think it's as much of the technology finding the right use cases than it is just amazing technology in and of itself.

[00:40:50] KP: When you think through those routine tasks that are going to be automated, maybe even your thoughts on the low-hanging fruit, are there any particular jobs you think should be most concerned about up-scaling? Whose heads are on the chopping block in the next five to ten?

[00:41:04] DT: Yeah, jobs that rely heavily on routine tasks are the ones most at risk. So that's everything from a stock broker. We've already seen the disintermediation of that industry. When it comes to robotics, we tend to think about the 3Ds, jobs that are dull, dirty or dangerous are good candidates for being replaced by robots. Real estate agents, things that can be highly automated via digital means, things that are highly replicable, I'd say those are ones where I wouldn't want my kids to be studying to go into those fields particularly in light of all of the amazing new fields that are getting created through the application of AI to traditional fields. There are certain ones that I think have a limited shelf life moving forward.

[00:41:55] KP: Absolutely. Well, the future is notoriously difficult to predict. One of the places I think people can get some good vision into where it's going, as your podcast that we mentioned at the top of this, AI and the Future of Work. Can you give me a deeper dive? What type of content do you get into?

[00:42:12] DT: Yeah, Kyle, it's my passion project. So every week I get an opportunity to sit down with oftentimes some of my heroes and have conversations like the one we're having about the future of humanity, the future machines, kind of this harmonious intersection of human and machine. We try to keep the conversations positive and actionable. So anyone who's looking at investing in that next career or the re-skilling, the up-skilling conversations. We look at a very technical pragmatic operational level at what are some of the trends? What are some of the technologies that underlie those trends? So that rather than a whole generation fearing the bot apocalypse, I feel like if we can do a little bit of good by helping fewer people fear the bots and more people kind of embrace the future of these amazing technologies, then we will have put a little debt in the universe, and that's what we try to do on the podcast.

[00:43:10] KP: Any particular technologies you're most excited about following over the next couple years?

[00:43:14] DT: I am. So I'm inspired every day by what our customers do with PeopleReign. And that kind of future of work vision that I shared about ambient intelligence, that's the one that we're sprinting toward. So we're super proud of the success that our customers have. We're super proud when our customers come back and tell us they're getting back 30, 45, 60 minutes a day, thanks to PeopleReign. But it really just makes me more ambitious about getting to that future of work vision faster and all the technologies related to things like taking IoT sensor data and incorporating that into machine learning models, or embedding sensors in traditional kinds of technologies that help us capture signals more effectively, or making it easier to integrate some of the systems that I mentioned previously into kind of one system of intelligence. These are all the technologies that our R&D teams are working on every day to be able to accelerate that vision for the future of work that we have.

We're millions of employees into this journey at PeopleReign, and it's a big planet. At some point I won't be able to say this, but today I can say most of the planet hasn't heard of or been exposed to the amazing technology that is PeopleReign. So I would just ask not as a vendor, but just as someone who cares about the future of humanity, technology and achieving this vision for a better employee experience. I'd encourage you, listen to the podcast, AI and the Future of Work. Check out PeopleReign. Go to our YouTube channel.

There's a bunch of recorded videos on there. Type PeopleReign into YouTube and you'll get a little sense of why I'm so passionate and so enthusiastic about what's ahead. And I think that the net impact of smarter machines is going to be making smarter humans. And I continue to just be incredibly motivated by what our customers tell us they're doing with PeopleReign. And we just keep pushing ourselves harder to make sure that we can deliver more value and really want to take pride in having played a little part in making all humans around the globe the best versions of humans that they can be.

[00:45:26] KP: Well, it's an exciting vision. Dan, thank you so much for taking the time to come on Software Engineering Daily.

[00:45:32] DT: Kyle, my pleasure. Thanks for having me.

[END]