

# AI and You

Transcript

Guest: Amit Gupta

Episode 64

First Aired: Monday, September 6, 2021

Welcome to episode 64! Just a word before we start: at the moment, it's August 2021 for me, now, and coronavirus is far from a solved problem. And I know that most of you are still struggling with it in one way or another. So first of all, thank you for having the spare cycles to listen to *AI and You*, because with COVID and climate change and so many other things it would be easy to get maxed out on global threats. And we look at AI on the show through many lenses, only one of which is "threat," where it doesn't have the immediacy of the things I just listed. But it does also have a relevance to both of those issues – COVID and climate change – through the vector of *disinformation*, which is something that spreads with the help of AI. And I mentioned that way back in episode 1 but haven't done much about it since then. So in a few episodes, I will. Get ready to pin back your ears.

Our guest today sent me the following bio: "Amit Gupta is an optimist, a science fiction writer, and founder of Sudowrite, the AI-powered creative writing app. He is also an uncle, a son, and a friend to all dogs. He previously founded and sold Photojojo and The Daily Jolt, and started a worldwide co-working movement called Jelly."

You'll want to get an idea of what Sudowrite is, so go to [sudowrite.com](https://sudowrite.com), and/or read the article in the *New Yorker*. Google will take you there. To call it a tool for writers is like calling a SpaceX Starship rocket a mode of transportation. Sudowrite and tools like it could take the very art of writing to the next level. If you've ever written a challenging piece, especially fiction, you know what writer's block is like. Your brain is just stuck like a sheep caught in a fence and all it can do is go baaa over and over. If you have a coauthor they can help but coauthors come with a lot of baggage, right? That doesn't often work out. Well, Sudowrite can fill that gap. It's an AI that will take a writing prompt and run with it, generating utterly plausible and wildly creative prose without ever duplicating anything already on the Internet. If nothing else, it blows open the whole school writing assignment homework thing, because no plagiarism detector is going to flag it. Give it the beginning of your essay, some context, and bingo. Without much effort, it will even sound like your own writing voice. I know we've talked about other uses of GPT-3 to do this kind of thing before, but you really have to see what Sudowrite does with it. You can even point it at a character in your story and ask it to create a description of them, and it will do it broken down into each sensory modality. Maybe you didn't think yet about how that CEO you wrote in chapter one smells, but Sudowrite will. And, once again, I did not use it to write this opening. I'll tell you when I do. Let's get into the interview now.

Amit, welcome to the show. It's a pleasure to have you on.

Thanks for having me, Peter.

And so you have been around since you had abundant energy in 1999, to get into the dot com boom at the peak, am I right?

Yeah, the peak and then the crash, I got to experience them both.

The crash. Well, if you're going to experience the crash, might as well experience the peak. What did that teach you?

Well, I was a snot-nosed, cocky, 19-year-old and I dropped out of school and had started this small company out of my dorm room and raised a bunch of money - or what to me was a bunch of money, today it would probably seem like a pitiful sum, I suppose - and hired a lot of people, grew really fast. And then when the bubble burst and all of our advertisers fled, I think I was humbled at 19 to have to lay off half the company, most of whom were older than me, or all of them were older than me, and lay off half the company again, and then spend the next year kind of walking in this grueling road back to break-even. So that kind of brought me back to Earth, melted my wings.

Builds character, as they say. And is a great lesson for later on when you're dealing with some other snot-nosed 19-year-olds who haven't had that experience.

Sure.

And then you created a successful startup, Photojojo. What was the premise for that?

Sure. So that was kind of a fun hobby of mine turned into a business. And the idea was to help photographers have more fun. And we did that by inspiring them with fun tutorials and tips on things to do with their photos. And also eventually manufacturing and designing accessories for both cameras and camera phones.

And the reason I want to get into your personal history here is that you had a life-changing event - nearly a life-ending event - around 2011, I believe. And I want to see how that affected your trajectory in this business because we're going to talk about how you got on to that. So, can you give us the capsulation of what happened to you then?

Sure. So in 2011, I was living in San Francisco and running this company we just mentioned, Photojojo. And the company was bootstrapped. I hadn't raised any money for it, never intended to sell, it wasn't your typical Silicon Valley company. And I loved what I did. And I got sick. For about two weeks, I had kind of a fever and fatigue, and stayed home from work, which is unusual for me, for that entire period. So I went to see the doctor, took some blood tests, told me to go home. And the next day, he called and said I had leukemia. And he said I had a couple of weeks or a week, I don't remember now, before I was gone, and that I needed to start treatment right away. And I remember I got off the phone, called my dad, who's a doctor, discussed kind of specifics with him, packed an overnight bag. And, you know, went to the hospital, took the next plane to Connecticut, then spent the next six months of my life in and out of a hospital getting chemo. I don't mean to go into too much detail, but I guess I'll give you a little more color. What I had was acute myeloid leukemia, which is a very aggressive form of blood cancer. And the only cure - and I say "cure" because it only gives you 50/50 odds of survival - is a stem cell transplant. And a stem cell transplant requires a very specific matching donor. So there's registries in every country to have catalogs of people who might be matches. I didn't have a

match. So I was very lucky in that despite not having a match, I was in a really good hospital, I could get chemo, I could have 25% chance of living, which was, at the time, very good. What happened next was a bunch of my friends found out what was happening, and they start to spread the word, start to organize donor drives across the country and across the world. And eventually, we did find a matching donor. So when I got my stem cell transplant in February, a few months after that diagnosis, it gave me 50/50 odds of survival. I was still kind of broken and beaten down after spending several months in the hospital getting chemo kind of non-stop, while at the same time running my company remotely. And I could barely walk down the hall at that point, but I spent the next six months recovering. And when I got back to San Francisco, I asked myself if I wanted this life back, if I wanted to sit back down in that desk chair and get back to work. And the answer was, no, I didn't. And so what I ended up doing was spending the next year or so of my life packaging up the company, getting it ready for someone else to run it, finding a buyer - because I didn't want to just leave it. I loved what we had built, and I loved the people that I was working with - so I found it a new home. And then I went out and did the things that you do when you think you might only have a couple years left. I traveled the world. I got a motorcycle. I got a dog. I did all those things. And then what next?

And it's been more than a couple of years.

Now, yeah, it's been quite a few years.

I want to look at what your friends did for you there because that was like, crowdsourcing lifesaving--

Yeah, it was--

Amazing.

Yeah, it was a miracle. I owe my life to them. There's hundreds of people without whom I would not be on this planet today. There's a few people who even took time off their jobs to run these campaigns for me, to find actors to do PSAs, to get press in all the major outlets and everything. And there's no way I would be here without them.

Wow, and has it given you any ideas of how the Internet might be leveraged to produce more results in the healthcare sphere?

Well, what happened to me is certainly miraculous. And I think that without social media, too, it wouldn't have happened because that's how we spread the message. So you hear so much that's out there, so much hand-wringing about social media and the evils of it. And I've witnessed firsthand some of the positives that can come out of it. But certainly, it wouldn't see what happened to me as a model for how healthcare should go. I think that without those friends, and without a lot of people who were connected, my message would not have gone quite so far and wide. And I've been able to help some other people who have been in similar situations and have reached out, and given them advice and made connections so that they can also spread their message, but it's not scalable. We need more people to enter the bone marrow registry, we need a better system, certainly for a lot of different reasons.

So now to fast forward to recent times, somewhere along the way you discovered interest in AI. Where did that come from? And how did it progress?

Sure. So when I left Silicon Valley, I started to do a number of different things and I eventually found my way to writing science fiction, specifically optimistic science fiction. And that's when I met James, who later became my co-founder. We were in a writing group together with a bunch of other tech adjacent people who are now reading science fiction. And about a year ago, we started playing with GPT-3 from OpenAI, and realizing some of the interesting things that this could bring about for writers. And as we played with it and started to show it to authors that we knew and respected and saw their reactions, we realized that there was something really valuable here. So that's kind of how it got started.

And I've played with it a bit. We're talking about a tool called Sudowrite. And I'm guessing that the "sudo" comes from the Linux command?

Mm-hmm, yeah.

Which means "take on full root admin powers to write". Or is it coming from the XKCD cartoon "Sudo make me a sandwich"?

I'd say it's both.

And you got exposure in *The New Yorker*, which was an article where the author gave Sudowrite the first part of Coleridge's poem, which is actually the only part that Coleridge got through, *Xanadu*, and it filled it in. And he found the result to be entirely believably Coleridge and beautiful, and I would agree with that. Now, for the benefit of people still coming up on the curve on transformers, what's the basic mechanism for doing something as miraculous as that?

Yeah, so what I tell people when they ask, "How does this work?" is that we're building on the backbone. We're building on what OpenAI has built. They've created this enormous model, a transformer model, as you said, where they've taken the common crawl, which is a crawl of basically the entire public internet as of a couple years ago, I think, and they've fed it into this neural net that has tried to condense it down into a piece of knowledge that helps to predict what comes next. So if you give it a string of words, it'll take that string of words and try to tell you what's most likely to be the next word - or more accurately, token, but let's say word. And so we're exploiting this ability to take people's writing and to have Sudowrite, predict what should or could come next. Obviously, there's no right answer there. So there's different settings within the engine to choose how, maybe creative, the engine is being in terms of what it thinks should come next. So some of our features will explain that to give you several different variations of what could come next. So the idea is that often writing can be a very solitary act. You're writing something, you get stuck, you're banging your head against the keyboard, you don't know where to go next. And if Sudowrite can give you five different directions from there, none of them may be the right direction, none of them may generate text that you actually want to copy and paste into your work, but one of them is probably going to inspire you and give you an idea for what you could write. We can help you get unstuck that way.

And let's just dig into the mechanism of this a bit and find a sweet spot between it being a complete mystery and a bunch of equations. At a very mundane level, if I give it the same input twice in a row, it will give me different outputs. How do you do that?

Yeah, it's non-deterministic. So I think if you set the temperature lower, using the API, which is - I actually don't remember how OpenAI actually explains the temperature, but we use it as a way to go from factfulness to more creativity. So if a lower temperature will result in a higher confidence answer. So if you're asking a transformer engine to fill in the blank of what is the capital of the United States, for example, you don't want to creative answer, you want a factual answer. And you want the same answer every time. So a low temperature there should get you Washington, D.C. every single time. We're using a high temperature because our target is creative writing. We want different answers every time. We want kind of creative out-there answers.

So does a high temperature boil down to bigger random numbers somewhere?

You know, I don't know. I don't know exactly how it's implemented.

And you had something I haven't seen in that form before - a choice of genres, I believe that you could pick for the writing. Can you tell us what those are?

Yeah, we have a feature called Twist, where you can put in a summary of your story, and if you're looking for some kind of an interesting new angle, or different direction to take it, you can select a genre. So based on the genre, it'll access kind of the common tropes within that genre and give you ten ideas each time you run it for different directions that story could go. And the way that works is that we've distilled stories in each of those genres down to their kind of component pieces of a plot. So we use that as an example to help Sudowrite understand what we're looking for, and the kinds of twist that can occur within that genre. So with that model, we're kind of tuning GPT-3 to give a twist along the lines of that genre.

And how did you do that? Did you have to go over hundreds of stories and categorize them?

Yeah, we manually went through, read a lot of stories, read summaries of stories, wrote our own summaries of stories, and then picked out a variety of representative stories that had a diversity of different plots.

Okay, so the genre now becomes an axis that things can be measured on.

Sure.

Oh, okay. And then you had another feature I haven't seen before, you could click on a noun, and say, "describe this character." I got to use this for writing a piece for a coach's workshop that I was doing, so I taught it "describe a coach." And it gave a nice description of what a coach does, and the mentoring, the guidance, and that and it said what a coach looks like and described them in distinguished terms, and then it started talking about what they sound like, and then told me what they smell like... and then what they taste like, and it started getting weirder from there. And then it got off into some sort of Lovecraftian, dystopian thing, which I

didn't show them, although, I'll save for a special time. How did you do the character description part?

Yeah, well, not to give away too much of the secret sauce, but I think if I were doing it today, I'd probably be-- or if we were doing it today, we'd probably be using daVinci Instruct. I don't think that's actually how it's built currently, but OpenAI gives you several different models that you can access through their API. One of the newer ones is daVinci Instruct. So it's their higher-level model, but instead of just having it continue, you can kind of tell it what you want it to give you. So you could potentially say, "Tell me what a coach smells like?" or "Tell me what a coach looks like?" and it would actually try to answer that question. So that's the simplest way you could do it. We're doing a little bit more pre- and post-processing because we want to take advantage of the context around that word that you select. So I don't want to get too much into that.

Right.

But I think on a basic level, you can actually just ask to daVinci Instruct to just describe that word and it should do it.

To me, it plays with an interesting concept in that one of the things that comes out of people's mouths very rapidly when talking about the capability of AI is, "It doesn't understand this, it's just being a glorified parrot." But it does seem that we're able to chip away at the boundaries of what *looks* like understanding to create some reasonable facsimiles of part of it. And this could be interpreted as that. It's adding a dimension that's more than just the transformer operation of autocomplete on steroids - "Let's predict the next sentence or paragraph." I mean, I clicked on coach and although I know it doesn't understand that that means a person, it nevertheless gave me descriptions of a person because I guess that that was what it was closest to in semantical terms.

Yeah, I think, we assign things meaning when we think they have meaning. And I think on some level, the most basic organisms are finite state machines that take inputs and generate outputs, whether they're actions or bodily functions, or whatever. And a higher-level being might look at humans and say, "This is just a machine that's intaking food and social inputs and outputting speech and excrement." And they might be right; from their perspective that could actually be accurate. So we can look at something like GPT-3 and say, "Yeah, we're just giving it a bunch of motions coming out the other end, there's no intelligence there." How do we define intelligence? That's probably a discussion that I'm not capable of or qualified to have. But as long as the output is something that's useful enough for our work, I think there's something interesting there.

So what was your motivation in creating this? It's a bunch of work with that categorization, and so forth. It costs money. And, was it to create something to make those lonely nights creating science fiction easier? Or where do you aim to take this?

Yeah, that's definitely the original motivation. I think James and I were using it as a tool in our own writings and we started to share with other people and they enjoyed it so we thought, "Okay, let's see if we can spread it more broadly." I think selfishly James and I also enjoy building this. It's a fun project and it's a fun sideline to the actual writing. When the writing becomes frustrating, we can make an excuse to work on the writing tool instead. It's kind of the classic programmer thing to do. But also, yeah, I mean, on a higher level, I think the future is going to show us that tools like this are very valuable. I think that other creative fields have certainly had tools of a lot of sophistication to help them create their arts, whether it's music or photography, or what have you. Many of these fields you can't even imagine doing today without software tools to create the sound or the picture that you're looking for. And yet writing has kind of been stuck, not that far removed from pen and paper. So we think there's a huge opportunity to change how the craft of writing works. And of course, as writers, we want to be involved in that.

So are you commercializing it?

Yeah, it's a paid product today. It's not a hugely profitable product. It's probably about breaking even at the moment. But our hope is that it can be something that's self-sustaining, and potentially also pay people to work on it. Hopefully this year or next.

What's the main cost? Is it cloud cycles now?

Yeah, at the moment, it's cloud computing. And I mean, salaries too, except there's no salaries being paid. It's just James and I at the moment.

So where do you see this going? Do you have plans for Sudowrite to [be] bigger and better?

James and I have talked about potentially building tools for other creative uses as well beyond writing, so that's in the back of our heads. But at the moment, I think we're mostly devoted to writing and to Sudowrite itself. We're taking directions from our authors. So there's a tension here, which you identified already, probably came up in the *New Yorker* article when you read it, which is, a lot of people are worried about AI tools taking over their jobs and making them irrelevant. And I think that our goal is to make something that enables humans and augments humans, not replaces them. And that's what Sudowrite is today. Or that's in its most charitable moment, what Sudowrite is today. I think often it can augment and help you, sometimes, it can just be a confusing mess. But our goal is to build a product that can act as a brainstorming buddy or a creative partner in your work. So that's the ultimate direction. How can we build something that makes the craft of writing less solitary, more creative, more playful, and just more fun?

A few episodes ago, we had a lawyer, Ryan Abbott, on and he was talking about intellectual property and ownership rights with respect to AI. Has that discussion happened for you with respect to Sudowrite, and what it outputs?

Yes, definitely. It's often one of the first questions that authors ask us when they can see what this does. Well, I guess, first and foremost, we don't take any copyright over anything that Sudowrite generates. So if you're using it to make your work, it's your work. And then the

second question is, “Is this actually original work that’s coming out of Sudowrite, or is it a regurgitation of someone else’s work?” And every time we’ve tested, every time we’ve looked at a particularly interesting string of words that is so poetic, it must have been written by a human, we haven’t been able to find traces of it elsewhere, haven’t been able to find it in plagiarism detection software. So we’re fairly confident, although we tell everyone, “Make this your own. This is something that we consider a way to generate a very helpful first draft, but not a final publishable draft. So make the work your own.” But yeah, it’s definitely in the back of our heads. I think a final product probably will incorporate some sort of plagiarism filter. And I think it’s a developing area for all these tools. How do we make sure that this work is unique and copyright-free?

What I think is fascinating and mind-blowing, is that the essence of the technology is that it’s a plagiarism generator. It’s just that it’s so good at it, that it doesn’t match our definition of plagiarism. But it’s taking things that other people have written and transforming them in such a way that it becomes something that no one has yet written.

Yeah, and then in some ways, that’s exactly what we do, right? We educate ourselves by reading all the things other people have written. And then we try to generate something that no one else has written.

Yeah. Right. And then, of course, there are all those times that people write something thinking, “That’s original. I’ve never heard of that before,” and it was unconsciously regurgitating something that someone else had done. And then sometimes they get taken to court for that and they go, “I didn’t know, it was an accident.”

Yeah.

One of the other functions of GPT-3 is you can run it in this adventure game-playing mode and pick a genre - say, fantasy - and then you’re in something like Zork. But it makes it up on the fly from what it knows and that’s a dialogue. What sort of memory does it have? How big is its context for what has happened to maintain the narrative thread? And you can answer in terms of Sudowrite. How far can it maintain a narrative thread before it forgets what it was talking about?

Yeah, so we’re using the full token window that’s available to us. So it’s the same as GPT-3’s, which is 2,000 tokens, or I think it’s 2,000 tokens, I may be wrong. Which basically, let’s say that’s like 1,500 words or something, and that includes both input and output. So there’s tricks for working with that. If you have a larger work, some of our tools will summarize the work and then take a smaller window of actual text to work from. So there’s ways to kind of massage it such that it preserves more of a window. And that’s something that we’re working on over time as well. As people create longer and longer works with Sudowrite, we want to find ways to preserve context, even when they’re much, much later in the work itself.

So you could extend that by taking out, say, the conjunctions and pronouns from the piece?

I haven't tried, I don't know if that would work. But I think there are certainly ways to condense it. The same way you might take a novel and kind of write a quick summary of each chapter, you can probably do something similar here to give the engine a summary of what's happened before and then tell it what's happening now.

Is it just English?

No. Although we're primarily working with English, some of our authors are working with other languages.

And did you have to do anything special to enable that?

Now GPT-3, since it uses common crawls, it's read multiple languages all across the internet. So whatever it understands about English, it understands to some degree about these other languages too.

Does that mean that if I go in and start writing in French, it will keep going in French?

Yep.

Can I confuse it by writing an English sentence that has a lot of French words in it?

I've written a short story where I had some dialogue in Hindi written in English letters, and it was actually able to continue that dialogue, but still have the scaffolding in English, so it can pick it up.

Wow. Has that or anything else that it's done surprised you?

I mean, just the fact that it works at all is kind of astonishing. Like the fact that it can continue a narrative thread, continue a character arc, and not on uncommon occasions generate a piece of dialogue or description that feels like art. Yeah, it can be quite surprising, in a lot of different ways.

It seems like ground zero of the AI revolution right now. You've got GPT-3, which can turn heads just by mentioning the name. You've got it generating text, it's flirting with the Turing test, it's right on the cusp of the robot revolution here. Is that a playground that you see yourself continuing in? Does that have an intrinsic appeal to you?

I think so. I think it has an appeal; plus, Kevin Kelly says you'll be paid in the future based on how well you work with robots. And I think that's true, whether it's appealing or not, to some extent. But luckily, as humans, I think as a species, we're very adaptive. And we're good at dealing with change and embracing it and finding the ways to take advantage of it. And with any sea change like this, there's a lot of opportunity to take advantage of the new tools and figure out new ways of working, new products to create, new ways of taking your ideas and turning them into something people want. And so I think there's a lot of opportunity today to work with these tools to do those things.

Now, this is working with English prose, it seems to prefer fiction; at least your site is oriented around fiction, it uses the terminology of fiction. I gave it some nonfiction, and it did reasonably well, but I get the sense that it wants to do fiction. Can that be extended to other things? It's language, it's just filling in language, it's filling in tokens. Those don't have to be prose in unstructured human language. They could be other things, I don't know, like, proteins, maybe. Is that a possibility?

Sure. Probably not with GPT-3, since it hasn't been trained on protein strings. But I think yeah, with any other piece of data, as long as you've got enough data as a training set, you could create a model that will spit out the results for sure.

Have you done anything with images?

We haven't, no. We're working with some new models right now to work on pacing, but that's again, English language words. We're really focused on writing at the moment.

Okay, and where is the next direction for you to take the writing?

There's a lot of stuff we're exploring. I think we have the luxury of, again, having raised no money, we can do the things that interest us and follow the threads that seem most promising to us as individuals, but we're also listening to the authors that are using the platform. So, one of the things we're starting to experiment with is a feedback tool. When I was building my company, I got very used to having very rapid feedback cycles. As a technologist, you can create a new product, you can A/B test something very quickly, and know within minutes if it's working or not, or hours. With writing, I learned the hard way that you can write a story and it might take you weeks or months to get a substantive amount of feedback on it. And that's a very long feedback cycle. So if we can create a software tool that gives you some feedback on what you've written, it gives you some ideas for things to try nearly instantaneously, I think it would really shorten that feedback cycle in interesting ways. So we've been building something along those lines.

When you say feedback, does it render an opinion in some way? Or is it you talking about how it continues what you've written and determining or asking whether that's a useful direction?

So right now, it will read your story and it'll try to identify the themes, it'll try to present you with some of the things that happened that it enjoyed, and maybe show you some of the things that it found confusing or didn't believe. And in the future, we're working on something where it'll also suggest other stories that are relevant to the one that you've written that you might want to read to explore further the ideas you're writing about.

I missed that mode. I was focused on the one where it was completing a story. So it can critique writing? There's a mode where you can give it something and it will tell you how good it is in different ways?

I don't know that it will tell you how good or bad it is, but it'll try to give you constructive feedback. And again, as with any feedback, I think the meaning of that feedback is in the eyes of

the beholder so we try to give you a few different versions. I belong to a couple different critique groups. And so I usually do a short story critique at least once a week for somebody. And lately, I've been doing that critique, and then I run their story through Sudowrite and tell them what Sudowrite said about their story. And almost every time, people have written back with a lot of surprise because Sudowrite identified a thematic element or something that they hadn't seen, that they found really interesting and decided to noodle with more. And like I said before, we find meaning everywhere as humans. So this thing could be saying things that are relevant to the story that maybe they're true, and maybe they're not. But either way, it might spark an idea for you that actually is going to be helpful.

In picking this audience of creative writers, it's a field where most of them are around the poverty level in terms of income. It's one of those things that's skewed. And are you looking at making up some of the revenue through places where there is money available for more automated writing, like annual reports, and ad copy, and things like that, where it doesn't require as much brain-bending creativity?

Yeah, it's a good idea. Most of the people who are using GPT-3 to create writing tools right now are focused on nonfiction, and they're focused on ad copy, and they're focused on website copy, that sort of thing. I think you're right that most of the potential dollars are there. It's just not something that we're that interested in so we're not really looking at it. And I don't know that we will, probably to the potential detriment of this as a business. But we think what's interesting to us is telling better stories, and there's a lot of value in creating stories. And there are certainly people that earn a lot of money telling stories, whether it's published authors who have a very successful series, or people who write for television or movies, or people who just write for joy and are willing to invest in a tool that helps them do that. So those are the people that we're targeting.

And thank you for giving that perspective because in this business, being so superheated, virtually everyone is focused on "How can I make money off the newest thing in AI?", and you're clearly looking at "How can I use it to increase the level of joy in the world?" and the money part of it is not your primary focus. So that's refreshing to see that. Where should people go to find out more about Sudowrite and Amit and future plans?

Sure. Well, Sudowrite is easy, they can go to [sudowrite.com](https://sudowrite.com) and they can sign up for the beta. And if they mentioned they found out about it in this podcast, we'll bump them to the top and get them in early. And as for me, if they want to read some of my science fiction, they can find it on my website, [amitgupta.com](https://amitgupta.com). Yeah, and of course, I should also mention my co-founder, James, also a great writer. He's at [jamesyu.org](https://jamesyu.org).

And I saw that you've recently attended the Clarion West Writers Workshop. Is that right?

Well, I did get in, and then COVID happened so it was canceled last year, so I'm deferred. They decided to do an online version this year and I wanted to wait for the in-person version, so I'll do that next year.

Good. It was on my bucket list a long time ago, but other things intervened in the way of that. What advice do you have for anyone that would like to get into this field? What truths and lessons have you learned that were hard-won?

I mean, I think my lesson or my advice wouldn't be any different from any other field, which is to look for interesting problems that people aren't paying attention to, stuff that you have, ideally an intimate connection to, and can't get out of your mind. I think with all of this stuff that I've done, whether companies or projects, it's just things that I've had a lot of excitement for and interest in, and that keeps me awake thinking about it. If it's not doing those things, it's probably not something you want to spend the next five or ten years of your life doing, so don't do it.

That's a perfect place to wrap this up on. Amit Gupta, thanks very much for coming on the show.

Absolutely. Thank you, Peter.

That's the end of the interview. There's a link in the transcript to the [XKCD strip](#) I mentioned. Check out [amitgupta.com](#) and some of Amit's writing, and of course, Sudowrite.

Amit's LinkedIn profile describes him as an "all-around nice guy," and that's one of the more accurate descriptions I've seen in that slot. I'm really glad you got to meet him. I like the idea of being able to change what Amit called the temperature of the algorithm to up the creativity of the answer. So what do you think Sudowrite and its ilk are going to do for writing now?

In today's news ripped from the headlines about AI, back in 2019, New York designer and TEDx speaker Sebastian Errazuriz announced on his [Instagram account](#) that "I think it's important that architects are warned as soon as possible that 90 per cent of their jobs are at risk. It's almost impossible for you to compete" with algorithms he said, adding: "The thing is you're not that special." He posted an animation of a parametric tool that can generate plans automatically, and you can see as someone changes the dimensions of a living space that it adjusts the positions of the bathroom, the number of kitchen cabinets, adds an island, that sort of thing, many others, in real time. He counsels architects to instead become software developers.

"Go into tech," he said. "Understand that those same spatial capabilities can be used in more abstract ways to be able to coordinate giant systems and develop the systems of tomorrow, working with other types of engineers. Not the ones who will pour cement, but ones that will write code. Please make the switch now. Don't lose your job."

Now, is he right about this? He has experience as an architect, of course, but not as an AI developer and marketer, so we can't take this as a guaranteed prediction of the effect of AI on architecture. It could just be the reflex reaction to seeing what that fancy automated layout tool could do. My wife and I designed our house, quite a few years ago now, using 3-D CAD tools, and while we might have been more picky and idiosyncratic than this tool would like for a customer – 'cuz I'm a control freak, right? - there's no question that it would bring joy to the heart of any housing estate or subdivision developer. Any of them could, by themselves, get satisfactory plans within minutes that would normally take days of back-and-forth with an architect. Whether that will put 90% of architects out of business, I can't say. I present it to you as part of the narrative that AI is creating for us, because accurate or not, Sebastian's

predictions are now part of the environment that architects are exposed to, the water they swim in; and what does that do to their thinking? Or the thinking of students in architectural school wondering how they'll break into the business and make a living? Even if we don't know what to *do* about the issue, we *can* look at how people *relate* to the issue. That's a very coach-y sort of thing to say, that's coming from my experience as a coach.

Next week, we are back in neuroscience, baby, in a massive interview with neuroscientist and TEDx speaker Professor Olav Krigolson, who figures out things about people from their brainwaves. He's used them to tell whether doctors are too tired to operate and whether Mars astronauts are ready for a mission, and he will tell you all kinds of things that you can do right now with off-the-shelf EEG hardware, and no, I didn't know there was such a thing as off-the-shelf EEG hardware until Olav told me about it. AI is intimately bound up with the way the human brain works because maybe the way to get general AI is to copy the one we know about already, which is sitting on top of your neck. But first, we gotta figure out how it works. We'll get into that 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>