

**SDS PODCAST
EPISODE 943:
CREATIVE
MACHINES: AI IN
MUSIC AND ART,
WITH PROF. MAYA
ACKERMAN**



Jon Krohn : 00:00 Algorithms compose music. So convincing that humans can't tell the difference between Bach, Mozart and Machine. And that has been the case since the 1980s. Why are we only now freaking out about AI creativity? Welcome to episode number 943 of the SuperDataScience Podcast. I'm your host, Jon Krohn. My guest today is Professor Maya Ackerman, associate professor of Computer Science and Engineering at Santa Clara University, where she specializes in gen AI research, a field she's been immersed in for decades. Today, we discuss her brand new book, creative Machines, and the profound implications of AI models that exceed human capabilities on more and more creative tasks every day. Enjoy this one. This episode of Super Data Science is made possible by Anthropic, Dell, Intel Gurobi, and Airia.

00:46 Maya, welcome to the SuperDataScience Podcast. It's a treat to have you on the show. Where are you calling it from today?

Maya Ackerman: 00:55 I'm in the heart of San Francisco.

Jon Krohn : 00:57 You recently moved there. I understand because your husband was a guest on the show recently in episode number 927, and I believe when we recorded that you guys were packing up to move into San Francisco proper and now you've done it.

Maya Ackerman: 01:13 Yeah, I'm becoming very good friends with the Bay Bridge here. It's a whole thing.

Jon Krohn : 01:18 Yeah, you have a great view at the window. I've seen the LinkedIn posts about it, of you improvising at your piano with a beautiful view of the bridge. It's fantastic. And so David, his episode number 927, it's one of my favorites that we've ever done on the show. He is unbelievably talented at explaining complex concepts in an interesting way, and I guess we should not be surprised then that he

also is very close to other outstandingly interesting people. And so I became aware of your work because at the end of his episode, he recommended your new book, creative Machines, AI Art and Us, which came out on October 14th in the US. It is already selling very rapidly. It has amazing reviews and yeah, sorry, I've been talking way too long, but I'm kind of just giving a bit of an intro to who you are. You're also, in addition to writing this recent book, you're an associate professor of computer science and engineering at Santa Clara University. You're a pioneer in the gen AI industry. You're a musician so much, but I think we should start with your book. What do you think about that?

Maya Ackerman: 02:35

Sure. That sounds great. And thanks for the kind words about my spouse that's always appreciated. Yeah, the book has been sort of accumulation of my life. I was so fortunate to enter Gen AI exactly 10 years before it became popular. So to have a chance to show a perspective that maybe goes against the grain of what we see in the press of how big industry, how investors see it, and paint a vision for the future that's much more human centric, but still AI forward.

Jon Krohn : 03:18

Yeah, I like that a lot. Your book talks about cutting through the hype, revealing the true capabilities and limitations of gen AI while championing its potential, while championing its potential to amplify human creativity rather than replace it. And I think you're right. I think there's a lot about replacement in the press, and so in creative machines, AI art in US, your book in the introduction, it warns us that this is a pivotal moment and it closes with a powerful call. You say that this book is a vessel to decide what kind of future we want to build. So Maya, in an ideal future scenario, what does a healthy creative ecosystem look like when both human and machine imaginations are shaping the cultural landscape?

Maya Ackerman: 04:06

Oh, the opportunity here is just incredible. We have these creative minds now living amongst us, these machines that we made, and sometimes we take this sci-fi led belief that we are creating our own replacement, that these machines become kind of the next evolutionary step and that we are going to be sidelined, go grow vegetables, enjoy your simple life is actually a relatively positive narrative today, which is absurd because what gets lost is that we are actually way more capable than machines are today, and we are nowhere near reaching the top of human intellect or the top of human creativity. And so us and machines can both move forward in tandem, but for that we need to design machines that are in their very nature aimed towards elevating us.

Jon Krohn : 05:09

So this human machine interaction, it is interesting. It is hard for me to understand where things are going. You have much better perspective than me. In chapter five of your book. You challenge Anthro centricism the belief that humans are the pinnacle of creativity. And you then in that chapter, you show that many of the prized traits that humans pride ourselves on exist throughout the natural world already. And so if we accept that creativity is not uniquely human, how does this change how we perceive ourselves? Are there implications for education, for ethics, maybe even spiritual views of innovation?

Maya Ackerman: 05:54

The risk is to insist that we are the only creative species or the only conscious thing or the only intelligent beings at the expense of actually figuring out how we should be living our lives with the presence of these machines. So it makes us uncomfortable that AI is creative. I get that. I can empathize with that. And what you find then is this kind of really wild oscillation between machines are going to completely replace us. We're not going to need any human musicians or any artists all the way to the other extreme. These machines are not really creative kind of denial. Instead, what we need to do is sort of admit where

we're at. If you look at some systems like Midjourney, like Suno, the fact that machines are creative is undeniable. We've had something called the discrimination test for decades. There was a guy by the name of David Cope who built a system named Amy, which made music and the style of Bach, Vivaldi and all the greats, and people would love the music until they hear that it's made by a machine, and then suddenly they could tell all along that it was machine made.

07:06 So they made a discrimination test where people would listen to the music and have to decide what's Bach and what's machine made, and they couldn't. I love that. It was so far past that.

Jon Krohn : 07:16 Roughly how long ago was that process created? That algorithm created,

Maya Ackerman: 07:21 That was 1980s

Jon Krohn : 07:24 Really. We had gen, ai, Bach, and Vivaldi in the eighties. The

Maya Ackerman: 07:28 Eighties. This is

Jon Krohn : 07:29 So old. Wow. I thought at most it was going to be like 10 years ago.

Maya Ackerman: 07:33 No, 10 years ago is when industry started getting involved. So we got things like Deep Dream, these cool hallucinatory images from Google, and we got things like Watson Chef from IBM, and that's when the press started publishing stuff about it. But in academia, we've had so much more time to think about it and develop it.

Jon Krohn : 07:54 So this algorithm from the eighties, what was it called again? Who was behind it?



- Maya Ackerman: 07:58 So this is David Cope. He actually, unfortunately just recently passed away. He was a professor at uc, Santa Cruz, and it was called Experiment in Musical Intelligence. That's what Amy stands for, and it's really, really good music. You can listen to it on YouTube. It's phenomenal. There's this great piece, Zodiac, a tourist piece that I just love competes really well with modern generated music.
- Jon Krohn : 08:25 Wow, that's cool. And so was it the musical notes that were generated and then an orchestra would perform it or?
- Maya Ackerman: 08:34 Yeah, so this was compositional generation, so this is when we're generating music sheets essentially the recent wave, the stuff that we see with Sun and Uio generates sound directly, which is a really, really meaningful step forward, but not as enormous as what the general public believes.
- Jon Krohn : 08:54 Right, right, right, right. Cool. Yeah, I would've been completely blown away if it was generating music that sounded convincingly. If it was generating the sounds in the eighties as well, that would've completely blown my mind. There are some pretty remarkable things that we were doing in the eighties with machine learning. I've seen videos of self-driving cars using machine vision systems using early neural networks from the eighties. Sometimes I do get really surprised, but I'm glad I'm relieved to hear that it was, I mean, it's still extraordinary that it was creating that composition. But yeah, it's kind of just in the last couple of years that I felt like convincing quality music generation has been possible. But maybe what was the earliest that you heard algorithms generating music? Not just the composition, but the sounds themselves and that it felt compelling, that it felt authentic. Authentic is word that



- Maya Ackerman: 09:53 Gave to Yeah, that's right. That was 20 22, 20 22. Some of the earliest
- Jon Krohn : 09:58 Work in
- Maya Ackerman: 09:58 This direction, the difference right now, the big leap is size and money, so both in music but also with large language models. The big epiphany was that epiphany in quotes is that in order for a brain to be smart, it needs to be big. We knew this long time ago in animals, the reason our brains are so good is because they're pretty darn big. We have a ton of neurons between our skull and OpenAI managed to convince Microsoft to give them an unprecedented investment to build a bigger AI brain. And so that's the recent shift, not the ideas, not the concept, not the generation. We've had plenty of amazing stuff. It's the money to make the brains bigger, and it's the money to market. So you can penetrate this sort of common consciousness.
- Jon Krohn : 10:53 And so when you say brains, it's like the number of parameters, the number of model weights in a large language model,
- Maya Ackerman: 10:58 It's like literal neurons. There's connections between them. We're building brains thinking about regulations. I've been thinking about that a lot recently. We're building brains. We are not able to tell them exactly what to do. Same as when we make a child, we don't get full control. So it's this new territory that we're really, really not used to the society with automation of we're creating brains. So the control is quite limited.
- Jon Krohn : 11:28 Yeah, approximations of cortical connections, the outermost gray matter of our brains that emulating the cognition that we're capable of, that humans are capable of. Yeah. I mentioned a word authenticity or authentic, and that's a bit of a loaded word actually. You have an

article we picked up on in our research where you argue that creativity should be assessed by outcomes rather than process, and I think that is supported by something you already said in this podcast where you talked about how human evaluators, even back in the eighties would love listening to these AI compositions of Bach and Vivaldi until they discovered that it was AI generated and then they knew it all along in your words. So yeah, you've argued in an article that we should be assessing creativity by outcomes rather than process, and so that could mean whether a machine did it or not, it should be the outcome, not that we throw it aside because the machine did it. And so in this article you talk about how humans as well don't necessarily feel the things that they write about or the music that they create, that you don't need to be authentic to create powerful music in a particular genre. So yeah, you've inhaled deeply, so I'd love to hear what you have to say.

Maya Ackerman: 13:02

Yeah, it's so complicated. I mean, I get it. I really, really do. The more I work with musicians and artists, the more I get it in humans. Art is such a deep form of expression, at least at its best. Often when an artist makes a piece of art or composes a piece of music, most of the time there is authentic emotion behind it. There is a real human experience to be shared. That's a big reason why we love art. Now, when it comes to professional art, that's not always the case. If somebody's commissioned to create this soundtrack and they only have, I don't know, 12 hours to do it or even two days to do it, they're going to get it done. Whether they feel it or not, it's their job. A photographer friend of mine has told me that he just makes up stories behind his photographs because the galleries demand that there be a story behind it. Ultimately, I think there is value in sort of expanding our perspective, at least in as far as being able to separate product from process. So if we are criticizing the process, we need to be aware that we're criticizing the process.

- Jon Krohn : 14:16 All right, very interesting. Another point that you've had that we pulled in from a recent interview is that you described large models as embodiments of Carl Young's collective consciousness, where we give the models permission to be as imaginative as possible. There's no notion of right or wrong then, but at the same time you acknowledge that it's all bits and bites and prediction. So do we need to teach young creators to recognize when they're conversing with this Carl Jung, this Jungian collective mind versus sampling from it? Yeah, maybe I can leave the question right there.
- Maya Ackerman: 14:58 So many interesting things wrapped up into one question. I love it. The biggest issue that we have today is that we imagine the AI to be an all knowing oracle. It's not uncommon to hear the word God compared to LLMs today. It's crazy, and it all comes from science fiction, this AI that's so much smarter than us and so young people and not so young people will go to something like Chade and trust it implicitly. And that point, that sense of trust is precisely the problem. We don't have to denigrate the ai, but that doesn't mean that we should be worshiping it. If we have a healthy sense of distrust, some level of doubt about what it gives us, that's enough to undo a lot of misinformation damage. That's enough to sort of curtail a lot of the risks. If the AI is not necessarily right, then when you share it with it, your personal problems, you're not going to blindly take its advice, right? No human is perfect. Even the smartest people on our planet are not invaluable. This expectation that AI is supposed to be invaluable is completely absurd. And so it's very, very important that we start moving away from that.
- Jon Krohn : 16:22 That makes a lot of sense to think about how LLMs are trained on some corpus of data. They're fine tuned in some way, and that whatever the training data are, you could never be guaranteed that it represents truth. And actually, oh

- Maya Ackerman: 16:42 Yeah, go ahead. I mean, I feel like, sorry, if you don't mind. It's not that there is something flawed in the way it's designed, right? It's that fundamentally this idea of an all-knowing oracle is ridiculous. So it's not that being based on prediction is somehow weakness. It's not that being tied to data is it's very much how we learn. We learn from sensory input, we predict things constantly. We're constantly predicting, but regardless of the process, again, we need to separate process from results. Regardless of how it works, it's not ever going to be perfect.
- Jon Krohn : 17:20 Yeah, that makes sense. And so speaking of imperfections, a common complaint of people using LLMs or conversational chatbots is hallucinations, and those are often seen as a negative, but you've argued that hallucinations are the underlying mechanism at work in creativity, whether that's human creativity or AI creativity. So how can we reimagine hallucination as a creative force rather than a flaw?
- Maya Ackerman: 17:48 Okay, let's start at an extreme how text to image models. You type a little bit of text and you get this whole picture in a few seconds, and then you type a little bit more text and you get another picture. Jon, I'm sure you've played with these models, right?
- Jon Krohn : 18:03 Of course.
- Maya Ackerman: 18:04 Alright, so very, very impressive so far from human capability. I think that's what most of us believe. Well, when you look at humans and certain psychedelics, we can generate images in our minds and videos way faster, way more amazing than what these models can do. So that's a starting point and a kind of really good way to demonstrate that these machines have not beat us. We just have not reached our own potential. But more broadly, we hallucinate all the time. We are constantly

hallucinating. We think we're perceiving real reality, and yet we perceive it completely differently from each other. So by recognizing the common roots between us and the machines as hallucinatory beings, even if it's just little hallucinations most of the time, kind of unleash this whole idea of control and precision all the time and recognize how much of brains rely on imagination is really, really key to lean into what's currently going on.

Jon Krohn : 19:09 Well, that certainly is interesting, Maya, changing gears a little bit away from your book, which to kind of just remind people, creative machines, ai, art and us brand new book that they can order, and of course we'll have that in the show notes. In addition to your book, you have long been, as you mentioned at the top of the episode, you've long been an advocate of human-centered gen ai, in fact, your CEO and co-founder of a business called Wave ai, which offers a unique approach to AI-based musical assistance with tools that include Lyric Studio and Melody Studio used by millions of people around the globe. Do you want to tell us about Wave AI and these tools?

Maya Ackerman: 19:50 Sure, yeah. This was all born out of my own struggles with songwriting. A lot of people who learned music as kids, I was trained on playing other people's music. It's really odd when we talk, we always make stuff up from our heads, but when we play music for some reason we're supposed to play other people's music. And so in the very first research project that I did around it, which was with David, he implemented the prototype. I would give it some lyrics like, it's so fun to meet Jon today, and then there are different ways you can sing it. It's so fun to meet Jon today or it's so fun to meet Jon today, right? Just millions and millions of possibilities and my ability to write songs really blew up at that moment. I had so much more freedom just by having the machine come up with different ideas for these tiny little phrases. And we ended

up opening up Wave AI in late 2017 and many companies, our third product became the successful one. It's called Lyric Studio. It helps you write lyrics. It doesn't write it instead of you. And my favorite part about what we created is that it makes you into a better pen and paper songwriter. So it has a positive impact on a core creative capability that lasts beyond your use of the software.

- Jon Krohn : 21:19 I like that. And so if Lyric Studio is the main, that's kind of been the most successful product, maybe we should dig into that a little bit more. How does it walk us through a user journey? If a listener was to go to Wave ai, I mean, so you just go to wave ai net and then you can click on the Lyric Studio or you could go to Lyric Studio net directly and people can try it free. So tell us what that experience is like when our listener goes to that site and tries it for free.
- Maya Ackerman: 21:55 So you can go in and the biggest part of the screen is just a text box for you to write, but whenever you're stuck, which might be right at the beginning, maybe you've never written lyrics before, so you need help right away, or if you're stuck in the middle of the process, on the right hand side are suggestions using our own specialized LLM that we built from scratch for lyrics writing. And it gives you ideas based on your topics, based on your unique writing style. It's never about being right, it's never about giving you the correct lyrics. It's about helping expand your creative vision to consider other options. People find that even if they've never written lyrics before with these suggestions, they can quickly craft something, but the more you use it, the better you get because it sort of naturally fosters learning and independence.
- Jon Krohn : 22:47 Excellent. And yeah, it looks like there's a free tier, but then even the paid tiers, I mean they're quite reasonable

for a SaaS product ranging from \$3 US a month up to \$10 a month in the US depending on which package you choose, how there's this kind of gold mode where you have better context detection and an advanced metaphors at a slightly higher monthly rate. So that is really cool. And then tell us about Melody Studio as well. It sounded like you started off when you were singing the song that you sang to us earlier about meeting me that well-known hit. It sounded like you were actually, it sounded like you had the lyrics in mind already and then so maybe that's where Melody Studio comes in and it allows you to have melodies generated.

Maya Ackerman: 23:41

Exactly. It's sort of the other two big components of songwriting in addition to is melody creation and then also chords. So it gives you ideas for chord progressions and then it gives you different ideas for melodies. And just like lyrics, studio artists use it differently. Some of them rely on the product really heavily for every single line. Other people would get one little riff, one melody, and then they can go off and write the whole thing depending on where they are in their creative journey.

Jon Krohn : 24:09

Nice. I like it. And so then with Melody Studio, is it generating, it's generating kind of sheet music like we were talking about that we've had from the eighties with the Bach and the Vivaldi, now you can have kind of original music in the style that you desire, and can you literally pick basically any style you can be like, I want jazz, I want this to be rock, I want this to be hip hop. What kinds of constraints are there around what you can do?

Maya Ackerman: 24:38

It's pretty wide. We cover most common genres. It's very, very flexible. It was fascinating throughout the development of this system to figure out which musical components are most representative of the genre and where genres overlap. That's just a random fun fact.

When you automate something, you get to learn to discover new aspects about it.

- Jon Krohn : 25:01 Really interesting. And have you ever performed music that you have maybe to a large extent has had the lyrics automatically generated, the melody automatically generated? Have you tested that with audiences or do you have feedback on how people perceive this music?
- Maya Ackerman: 25:18 Of course. I mean, first of all, people at Universal Music use our systems.
- Jon Krohn : 25:24 There you go. So
- Maya Ackerman: 25:25 There is popular music that uses our stuff. There is a guy named Curtis King who had a number one hit on iTunes, whole hit album where he used lyrics studio. One of my favorite stories is actually super early on, my good friend James Morgan always wanted to write an Italian aria. The problem is that he didn't play any musical instruments and doesn't speak Italian. So we built a little system for him, trained on a public domain, works of Yoko Puccini. Anyways, after a couple of months with something that would just spit out tiny little vocal melodies, given Italian lyrics, the guy wrote this amazing piece called Arri ta Raho about this world of Warcraft character writing a dragon. It's beautiful Puccini like piece that really blew my mind about how far AI can help people go. And of course, I performed a ton of songs myself that I made with this over the years.
- Jon Krohn : 26:25 Nice. All right. That's cool. So that's amazing. Universal music, using it, these hits coming out. I've got to try this myself. Before we started recording, I was griping to you about how it sounds like maybe you were like this as well. I've been trained to play classical piano to play guitar, classical vocal performance as well, and almost entirely that has been about just regurgitating compositions that

have existed for centuries. And I never was taught how to be creative and how to create my own music. I've done a tiny bit of it, but I feel like it's a real weak point for me musically. And these tools, lyric Studio, melody studio, they sound perfect, especially the way that it sounds like you've calibrated these tools in a way to work where it doesn't feel like I'm being dictated at, but I'm actually working alongside these tools to create lyrics and melodies.

Maya Ackerman: 27:26

Exactly. Jon, I think there's a future famous composer inside of you.

Jon Krohn : 27:33

We'll see, we'll see. We'll have lots of great hits about all of your favorite data scientists from history. That sounds amazing. Have performed by me. Yes. Starting with Y Ko, he's going to be our first, I dunno, I'm making all up. So let's now kind of bridge both your startup as well as your book. So you've previously argued that the sci-fi dream of AI as, yeah, I mean we talked about earlier as an infallible answer, machine is disempowering and it's misguided. And in chapter 10 of your book, you note that disempowerment doesn't scale, but human flourishing does. And so how can investors and companies reframe their perception here? So reframe engagement metrics away from stickiness and towards sustained creative empowerment, kind of like you and David have with Wave ai.

Maya Ackerman: 28:38

Yeah, it's interesting that this addiction model persists even when the biggest problem with gen AI products is this kind of one hit wonder phenomenon. When a person feels really, they're really unnecessary. Let's say there's a little app where I upload my photos and then it does something fun with them, I might enjoy it once or twice, but if I feel like I'm not really doing anything, I don't have any control over what comes out, even if it's kind of cool, I'm very unlikely to come back. So we had a whole bunch

of one hit Wonder Gen AI products. People want to feel that they're doing something. People want to express themselves, they want to realize their own ideas. And industry and investors in particular have been very, very, very, very slow to realize that. I believe that the reason Chacha PT is successful is because for those who want to give of themselves, for those who want to really collaborate, it makes that possible and that's why it's so successful. So it's sort of like we need to snap out of these old outdated exploitative practices that a lot of us sort of believe in them as gospel and really explore what's meaningful and appropriate with this technology, which I believe from my experience in business and otherwise, that the real potential is to really elevate humans in our capacity.

Jon Krohn : 30:12 Yeah, it's a great idea and I agree with it wholeheartedly. It does seem like it's going to be hard to convince product managers. Do you have any sense, I realize this is a really tricky question. I don't have a good answer to this, but you're a lot more creative than I am. How do we convince big enterprises product managers to move away from just stickiness to empowering people?

Maya Ackerman: 30:38 I think we need to explain to them why something like why the products that are successful, why they are successful. ChatGPT has not been beat yet. It's the number one, it's one of the most successful products in history. And that's the case, not because you press a button and it replaces you, but because of power users, you always have to look at power users because those are the ones who are the customers that love you, show you where the product needs to go, and those are the users that go deep that become better as a result of using chat. I've heard people say that their vocabulary expanded. The writing style got more diverse as a result of using ChatGPT when you use it with a certain kind of intention. And also to show examples of how this sort of replace of

paradigm has often failed in a lot of smaller products, but also how it ultimately fails collectively. This whole, everyone optimizes for themselves really fails when we want to apply AI to replace human workers because if we don't have any more human workers, the entire economy collapses. And so it's really time to wake up from these incredibly greedy principles that are guiding our economy and to think more holistically in this moment, not assume that the old principles are just going to keep working and somehow magically everything is going to work out.

Jon Krohn : 32:03 Perhaps gen AI will kind of force this shift that you're hoping for in product managers and in enterprises. You mentioned in your most recent response this idea of diverse responses. And so in interviews you have previously contrasted convergent systems that gravitate towards safe average outputs with divergent co-creative tools, presumably like your lyric studio and melody studio that widen the search space. And like you just said, broaden people's vocabularies, increase the diversity of their writing styles. So as AI scales to creating music, to creating video, how do we keep models from nudging creators toward averaged risk averse aesthetics where you hear another pop song that sounds the same, how do we instead get diversity?

Maya Ackerman: 33:05 Yeah, it's really impressive what industry has done to gen ai. Gen AI has always been about expanding possibilities. That's how it was born. Going to new spaces, we think of creativity as this massive search of possibilities. The next line of lyrics, there are billions of options. How do we search, right? You're about to play the piano. Jon, you were talking about playing other people's pieces, like we're all told, where do we even start? What could be my first note? What could be my second note? It's a search space problem, and machines are phenomenal at just exploring different unlikely possibilities and even in trying

to take you to places that are pretty good, but not very common. The problem is that science fiction has convinced investors and some entrepreneurs that it's better to create an all-knowing Oracle that gives you the most expected, most safe, most.

34:04 You've seen it a million times already answer, which is maybe good for a fact lookup table, but it really limits what gen AI can do. And it's never actually going to be very good at being an all-knowing oracle regardless of what we do. And so when you look at chat chit and you're like, oh, AI is not creative, that's because chat chit is optimized for the opposite. Inherently, from a technical standpoint, you have to take creativity into account from the beginning in the way that you construct this machine brain. But in the grand scheme of things, in the grand scheme of things, it's not that hard. If we figured it out as a team of three, granted we have David, but still a team of three people originally at Wave ai, then I'm sure somebody like OpenAI, Microsoft and Google can figure it out. It's just not their goal. Their goal is not to open our minds. Their goal is to replace search. And I think that goal needs to be modified.

Jon Krohn : 35:03 I like that perspective. Hopefully we can get there, hopefully we can. We can iterate towards these divergent creative supportive systems that are allowing human creativity to flourish and allow humans to flourish as well. More broadly, let's talk now a bit about your academic work. So you're a professor at Santa Clara University. How much is the overlap in what you research? How much is there an overlap between your research? And what we've already been discussing in this episode is your research a lot about creativity. I see that a lot of your most cited papers, for example, are about clustering. And so I'd love to hear what the overlap is between your interest in creativity and what you're doing academically.



Maya Ackerman: 36:04

My PhD work and a little bit afterwards focused on foundations of cluster analysis, which was a much larger research community, and that's why you see more citations. It's really, really simple. And then I entered the niche world, which my department chair was not pleased with me for switching into Gen AI back in 2015, not what I was hired for. And so that's why it doesn't bubble up to the top quite as quickly. I've done all types of stuff in my academic work. I did a lot of work on bias in these models and sort of also a bit of a Jungian lens at that one in particular is called brilliance bias. So brilliance bias is this belief that we all have to some degree, which is that exceptional intellectual brilliance is a male trait. So basically when people think genius, they think male and they mostly think Einstein. And we were able to show that generative image models have the same bias. And so what does that mean? We're essentially perpetuating biases through the ai. We're just mirroring what people already believe and often amplifying it. But it's interesting, we can rush to kind of get mad at the companies, oh no, how dare you be biased. What were you thinking? But it's also reflecting our world, and I think in that sense it's very interesting. It's an opportunity to learn about ourselves a little more.

Jon Krohn : 37:37

Yeah, so I think I found that paper here. It looks like it was presented at IEEE at the Global Humanitarian Tech Conference, GHTC in 2022. And so I'll have the paper for people to read in in the show notes. It is kind of amusing to me a little bit meta of that. So you're testing at that time what were the cutting edge models in 2022 GPT-3 and the specific variant of GPT-3 that was most common and most sophisticated was called DaVinci. And so it's funny that you're studying this brilliance bias in a male trait and the models that are literally named after people like Leonardo DaVinci,

- Maya Ackerman: 38:22 The world can't help it because women did not have educational opportunities and whatnot. This idea of only men being true geniuses is easy for that to sustain. There is also a conference paper at the International Conference on Computational Creativity about this phenomenon in image models, and we're doing a journal paper right now with my PhD student, Juliana Shihade.
- Jon Krohn : 38:45 Really? Cool. We'll look out for those. And I'll include in the show notes all of these papers for sure. What else is really exciting? You academically, what's the Tara incognito that you're hoping to traverse into next?
- Maya Ackerman: 39:02 Oh my goodness. Where do I even start? You touched a little bit on the Jungian stuff, but I didn't properly zoom into that.
- Jon Krohn : 39:11 Yeah, let's do that because there are papers here that I'm seeing recent papers, for example, there's a 2024 paper called The Collective Mind Exploring our Shared Unconscious via ai. That's really interesting. And it looks like, if I'm understanding this correctly, in this Google Scholar citation, it looks like it's in a journal called Religion
- Maya Ackerman: 39:32 Or Religion. We'd have to verify that briefly.
- Jon Krohn : 39:38 I dunno how that happened. Must be. I think it could be a Google Scholar issue. No,
- Maya Ackerman: 39:44 Because
- Jon Krohn : 39:44 It looks like it's actually conference proceedings.
- Maya Ackerman: 39:46 Yeah, it's a conference proceeding. Yeah.
- Jon Krohn : 39:48 Wow. That is weird. What a weird thing for a Google scholar to do

Maya Ackerman: 39:52

Got creative on us. It did indeed. Can't trust anything. Indeed. So Carl Young, I find to be incredibly inspiring, a really, really deep thinker. I went back to his original writings and I felt like I got to know him a little bit. So he wasn't really the first, but he really expanded on the idea of a collective consciousness. So especially in our society, we like to think of ourselves as independent, just a whole bunch of people with their own unique brains and their own unique ideas. Now of course, that's complete nonsense. We tend to share fundamental beliefs about life with each other to an incredibly high degree, and each culture has sort of its own belief and its own belief system that its members subscribe to the very, very high degree. So what we think of as objective reality or as the right way to understand the world is often just being part of this collective consciousness.

40:46

So we talked about this a lot and how it impacts us and what that means. It's super fascinating and I encourage everybody to dig more into Carl Young's work, but what's really amazing, and I think Carl Young would've been so excited to live today because it's this collective consciousness come to life. We took all this western data and we created a brain out of it. So now if we want to know what Western consciousness thinks mod as alignment that all these companies are doing, we just talk to the LLM or go to something like Midjourney, which is not as aligned the text to image model, and it just reveals Western biases like there's no tomorrow. For example, I typed Hanukkah into Midjourney and I got a German Christmas tree. They fixed it since I published it, kind of revealing that people in the west tend to associate Hanukkah was Christmas, has nothing to do with it whatsoever. It's not even a very important Jewish holiday. I asked for T, which is a Hanukkah food, and I got bagels kind of associating Jewish food with bagels. Too much so beautifully honest about Western beliefs more so than any person would ever be. So I think there's just so much

to studying there beyond just yelling at the companies and we should get the companies to fix it, but in the meantime, we get to understand ourselves better.

- Jon Krohn : 42:13 That is interesting, and I have, it's some years ago now, but I went down a bit of a Carl Young rabbit hole myself, and he did have a lot of interesting ideas. So I can see how you got so into that. So we've covered in this episode already how you are an accomplished CEO book author, university professor, AI researcher. I've mentioned how you're a pianist and how there's lots of videos that I've seen of you improvising on you record improvising at the piano and have published that regularly on LinkedIn, so I've seen that, but our research also dug up that you're an opera singer.
- Maya Ackerman: 43:08 Yeah, that's right. I had a very unusual journey with music. I started playing piano when I was eight years old, and when I was 12, then my family moved to Canada. They were not able to take my piano with them, and so there was this massive 15 year gap in my music education. I got back into music as an adult when I was 27 years old. I started taking voice lessons, which was actually a callback to my childhood in Israel. I used to perform. I was recorded on national television as a kid, but then there was this massive gap, and so I learned how to sing opera in my late twenties and I started performing. My love of creativity is very, very real. I like to make music. I don't care if other things and machines can also make music. I want to sing. I want to play piano. I make art. I like to take classes that challenge me in traditional art forms and music. I miss opera performing. I should really do more.
- Jon Krohn : 44:12 I was going to ask if there's any way that we can check you out performing in opera. It sounds like maybe nothing is lined up at this time, but we can,



- Maya Ackerman: 44:20 But there's a lot of older stuff on YouTube.
- Jon Krohn : 44:23 Cool. All right. We'll try to dig some of that up for the show notes as well. So Maya, now looking forward, what are the kinds of things that our listeners should maybe be able to expect with the gen AI systems in the future? So obviously today we've become pretty good at music generation. It's been a couple of years now that we've had high quality image generation that all hands are getting five fingers and that kind of thing, and video generation has come a long way in recent months as well. Is anything, do you think you have any, maybe you've been in this industry for decades in a way, studying gen ai, so do you think you have any particular insights that we might not be expecting about what's coming next with Genai?
- Maya Ackerman: 45:12 There's a lot of focus on replacing different industries, different types of workers. I believe that's been the main effort of investment funds in the space. Pick a career, try to replace it with ai. But people are also learning that most of the time it works much better when you do have a human involved and they're slowly starting to recognize that. So with things like images we have in painting now where you can kind of highlight a section and regenerate it, SUNO also learned that its best users are people who actually want to make music. So they started to make that more interactive. So instead of just sort of more and more and better and better make these machine brains able to interact with others, you can think about really, really brilliant people who don't know how to interact with other people. So that's kind of what we have right now in AI to some degree. Now, what I hope we're going to say, and there is a movement in this direction, is that these brilliant AI brains and all these different industries are going to become more social, more able to collaborate with human beings.

- Jon Krohn : 46:23 Fascinating. I like that. Yeah, so I mean, I guess taking the kind of approach that you've had with your own wave AI products where we're more collaborative and it's less, it's not just a simple providing a prompt and getting a perfect output, it's about iterating and as you mentioned, being able to highlight maybe parts, very specific components of an image or a video or a piece of music, and being able to iterate on those particular elements as opposed to regenerating the whole thing from scratch and wondering where you've gone off to.
- Maya Ackerman: 47:01 Yeah, exactly. Isn't it frustrating, right? It's almost there, but it doesn't understand. When you say change the color of the hat, it can't.
- Jon Krohn : 47:10 Yeah, for sure. Actually, recently, I made a mistake in my notes for our production team. I basically, I copied the wrong file from one episode into another episode, and so we had a thumbnail, a YouTube thumbnail generated, created by our thumbnail designer with the wrong title. It was a title of another episode, and I thought, huh, you know what? I think I can use AI tools. I tried a couple, I tried Gemini, I tried chat GPT to upload my image and change the text to the correct title without changing anything else in the image. And the one that I happened to use first was Gemini. I'd heard that it was quite good at this kind of thing, and there was a spelling mistake in the generated text. It looked perfect, everything was perfect. It retained the imagery, but there was a spelling mistake. I had this word accelerated, and one of the E'S in accelerated was just missing.
- 48:18 It was correct in the prompt that I gave it. It was missing and it could not fix it. I went through half a dozen times of saying it kept saying, here you go. I fixed it and it just keeps printing out the exact same image with accelerating misspelt in the exact same way. And so that was a really interesting experience, and then I just moved over to chat

g pt, and it happened to get it right away. So yeah, it's an interesting, I think there's a lot of opportunity in having all of these generative tools work better with us still lots years of progress, hopefully to be made.

Maya Ackerman: 48:56

Yeah, it is also really interesting how it struggles with this sort of task, which we consider fairly straightforward, how it still struggles with some basic math, although it used to be even worse, because these are fundamentally imagination engines, and these companies had to put a lot of layers on top of it to try to make it behave, try to get it to go in a right reasonable direction while it keeps wanting to do its own thing, kind of like a kid.

Jon Krohn : 49:25

So for the math thing, for example, instead of using model weights to approximate numbers based on similar kinds of math, it's memorized from the internet to actually generate some Python code to do the calculation instead or something like that. Yeah, these layers, the layers on top, that's become a big part of the secret sauce of any of the frontier labs

Maya Ackerman: 49:46

For sure. For sure. Which makes sense. I think it is helpful for you to be able to ask it to add some straightforward text without mistakes. This is all valuable for sure.

Jon Krohn : 49:55

Yep, yep. Well, Maya, it's been a fascinating episode. Thank you for your rich perspective. I don't think we've had an episode like this on creativity before at all, and I don't think we've ever talked about music so much in an episode either. So thank you very much, Maya. Before I let my guests go, I always ask for a book recommendation, and so that's how I found out about you in the first place, was by asking David, your husband for a book recommendation. Maya, do you have a book recommendation for us?

- Maya Ackerman: 50:24 I'd like to recommend the work of a Seth who is a neuroscientist and professor, and he just does a phenomenal job explaining why we hallucinate all the time, developing a little bit of empathy. It could help us develop a little bit of empathy for the way that machines work as well, and not to imagine that they hallucinate and we don't, but also kind of reframe imagination in a more positive light. There are some wonderful YouTube videos that he did, but in particular, there's also a book called Being You a New Science of Consciousness.
- Jon Krohn : 51:00 Yes, yes. I found that as you were speaking about, it looks like a great recommendation. It seems like it ties together a number of themes from today's episode as well, so thank you so much, Maya. For people who want to get more of your insights or maybe watch you do outstanding improvised piano from your beautiful view at your office there on LinkedIn, where should people be following you?
- Maya Ackerman: 51:25 LinkedIn is a main platform that I use. You can easily find me Maya Ackerman on LinkedIn. Yeah, great to connect there.
- Jon Krohn : 51:34 Wonderful. Alright. Thank you so much for taking the time out of your no doubt, very busy schedule and yeah, maybe we will have to check in again in a few years and see how the gen AI music generation in particular industry is coming along.
- Maya Ackerman: 51:50 That sounds great. Thank you so much for having me, Jon.
- Jon Krohn : 51:56 What an interesting episode, and it Professor Maya Ackerman covered how we should assess AI creativity by outcomes rather than process. Where the machines are involved is beside the point. She talked about how her company wave AI built Lyric Studio and Melody Studio as co-creative tools that help millions of users write songs.



She talked about how large language models function as manifestations of Carl Young's concept of collective consciousness, and she provided her vision for the future of Gen AI wherein machines become more collaborative and better at understanding our creative requests. As always, you can get all the show notes including the transcript for this episode, the video recording, any materials mentioned on the show, the URLs for Maya's social media profiles, as well as my own at superdatascience.com/943. Thanks to everyone on the SuperDataScience podcast team, our podcast manager, Sonja Brajovic, media editor, Mario Pombo, partnerships manager, Natalie Ziajski, researcher Serg Masís, writer Dr. Zara Karschay, and our founder Kirill Eremenko. Thanks to all of them for producing another stellar episode for us today for enabling that super team to create this free podcast for you. Oh, so deeply grateful to our sponsors. You can support this show by checking out our sponsors links, which you can find in the show notes. And if you'd ever like to sponsor an episode yourself, head to jonkrohn.com/podcast to learn how to do that. Otherwise, support us by sharing this episode with folks that are looking for creative inspiration from machines. Review the episode on your favorite podcasting app or on YouTube wherever you consume the show. Subscribe if you're not a subscriber, but most importantly, just keep on tuning in. I'm so grateful to have you listening and hope I can continue to make episodes you love for years and years to come. Till next time, keep on rocking it out there and I'm looking forward to enjoying another round of the SuperDataScience Podcast with you very soon.

