

AI and You

Transcript

Guest: Kenneth Stanley and Joel Lehman, part 1

Episode 135

First Aired: Monday, January 16, 2023

Hello, and welcome to episode 135. My guests today are Kenneth Stanley and Joel Lehman, authors of a unique book in the field of AI, which, despite being published in 2015, I only just came across it, for reasons that are about to become apparent. The book is titled [*Why Greatness Cannot Be Planned: The Myth of the Objective*](#), and the cover sports a picture of a butterfly. The book's description starts:

Why does modern life revolve around objectives? From how science is funded, to improving how children are educated -- and nearly everything in-between -- our society has become obsessed with a seductive illusion: that greatness results from doggedly measuring improvement in the relentless pursuit of an ambitious goal.

So why are we even talking about what's clearly a management and leadership type of book? Because Stanley and Lehman are artificial intelligence researchers, and it was a discovery they made in the testing of the program that created the butterfly-like cover image that led to their writing this book, which for all its focus on inspiring leaders, is shelved under Computer Science. We'll hear why in the interview.

Ken Stanley was previously Charles Millican Professor of Computer Science at the University of Central Florida and was also a co-founder of Geometric Intelligence Inc.

Joel Lehman is a machine learning researcher interested in algorithmic creativity, AI safety, artificial life, and intersections of AI with philosophy and psychology.

Both Ken and Joel were at Uber AI Labs, where Ken was head of Core AI research and Joel was a founding member, and they were both again at OpenAI, co-leading the Open-Endedness team (studying algorithms that can innovate endlessly).

Let's get into the interview.

Peter: Ken and Joel, welcome to *AI and You*.

Ken: Thanks for having us.

Joel: Yeah, great to be here.

Peter: Okay, so thank you, and I want to talk about your book, *Why Greatness Cannot be Planned, the Myth of the Objective*. And what's striking to me most about this is that from the cover, there is no clue unless you're one of a thousand people that would know what the significance of the image is there, that this is about computer science, it looks like something from Marshall Goldsmith or Peter Drucker, it looks like it should be shelved on "self-help" or "business leadership" and yet, the back of the book Springer says, "put this on the computer

science bookshelf." And it could be a long time in reading the book before you get to the point where you realize that's where it belongs. Am I alone in this kind of perception? Or is there a story about how it came to be this genre spanning?

Ken: I think that you're not alone, probably. There is some story there. I mean, the research that led to the book comes from computer science and in fact, Joel and I were doing kind of hardcore AI, computer science research, and presenting in computer science circles, which are generally computer science conferences and journals. But what happened was that Joel and I, at some point, it kind of dawned on us, or just gradually dawned on us, that the lessons that we were learning from this work had much broader implications - like way broader - like to general, institutional, cultural, and even personal issues that go way outside of computer science. And so as we realized that, and we sort of decided together that, let's try to write a book that speaks to all of that, which is basically everybody. That's very different than computer science, the usual stuff that we were used to doing. But we thought that it was really important to get this word out more broadly, because we started to realize that the message is relevant much more broadly. But then you can see that from there, there's a kind of a dilemma, because it's like, well, what is the subject of this book? I mean, it's very heavily based on computer science research, but it has this kind of larger social critique, general cultural type of message and so it's true it was very hard to categorize. I mean, for us, even like to say, what is the genre of this book, and I think, and Joel should also chime in. But just because of my perception, sort of that we didn't try to actually pigeonhole it and just said, we'll just make it what it is sort of well, look, let's just disclose the research. It's really interesting research behind [it], and then make the point, that's very important, we think, and it's going to be hard to categorize it, but we're basically being honest to our message, and the chips will fall where they fall basically.

Peter: Sure, so maybe you could tell us in your own words, what that message is.

Ken: Okay. The message of the book, I think, is that when you set an ambitious objective, which is, by the way, ubiquitous across our culture, I mean, all of our institutions run on objectives. If you're applying for a grant they want to know what is your objective, if you're trying to do something in a large corporation, there's things called OKRs: objectives and key results and they want to know, what are the objectives or else there's basically no justification to do it at all. If you're doing something in education, you need to have some objective, like you're trying to get a certain test score or something like that to pass on to the next grade. Everything is objectively driven. And so what's interesting, though, in our message is that if you have an ambitious objective, and the word ambitious is important, because when it comes to modest objectives, I think this critique wouldn't apply. But when you have an ambitious objective, it can get in the way of achieving that objective; which is very counterintuitive. It's like by setting the objectives you're actually hurting your ability to achieve itself and it also gets in the way of doing anything interesting at all, on top of that. So it's basically very toxic to set ambitious objectives and this is, of course, not widely recognized, because if it was, it wouldn't be so pervasive across the culture and of course, it's counterintuitive, probably controversial in a lot of quarters, and so it's a hard thing to digest. But if you think about it, if it's true, which we spent a whole book trying to argue, it's super important because we're doing a lot of stuff wrong if this is true, where

everything is objectively driven. And of course, we speak a lot in the book about the alternative, because the first thing that comes to peoples' minds is, well, what am I supposed to do just act randomly? Like, how can I have a life without an objective? But actually, there are a lot of alternatives. I'm sure we'll get into them. So, that's the other thing: we try to provide the alternative, to explain what kind of world it would be without everything being objectively driven.

Peter: It's very Yoda: Do or do not, there is no try. At least I feel that kind of philosophical elf on my shoulder whispering in my ear when I'm reading this book. Joel, at what point in the computer science research that you were doing did the idea of the book start to form for you?

Joel: It came across both of us was gradually that abstraction from computer science, a search algorithm could have broader implications. It's yeah, I thought it was strange and also unexpected thing for that to happen and think Ken probably had more of the kind of realization of how strongly the message was resonating outside of the computer science sphere that would kind of prompt the possibility of launching a popular science book that's more generally aimed.

Peter: And was it during the investigation and study of the PICbreeder program that this came to? That's central to the book, was that the genesis for this idea?

Ken: Yeah, no PICbreeder was preceded. So, there's pictures like the first step, I think in a succession of insights that led to the book, because I think what really happened is it PICbreeder led Joel and I, to realize that there's another way of thinking about algorithms. So, it wasn't yet to the book is sort of at a social cultural level. But at the time of PICbreeder, we were like, well look like, we're seeing a phenomenon inside of this system, where people are only able to find interesting things when they're actually not trying to find those things. And this is very counterintuitive, and was really striking to me fascinating and I spoke to Joel about it, he had just started his PhD at that point. So, he had just joined my group, and we were just getting to know each other and I was like, just basically "Joel, this could be a huge thing. Like, look at this thing here. It's very provocative." But we were both I think, just thinking in terms of algorithms and so this led to the novelty search algorithm, which Joel basically took that and like, converted it into an actual thing, you could actually run on a computer without people involved. So, basically, Joel took the lesson and basically made it actually something that was an algorithmic contribution in its own right, which was amazing, I thought, I mean, this is like a really great, provocative algorithm. And it was after that, like the novelty search percolated around for years, people commenting on it started controversy in the field, because it was like, in the field of machine learning objectives are also completely pervasive, just like in the wider culture. Algorithms in machine learning, you basically say, "we'll set an objective and then we'll climb a gradient to try to get to the objective" and it's basically like the way algorithms work. And the novelty search algorithm was basically, throw all that out. It's actually bad for you, the algorithm won't even know what it's trying to do, but actually do better in some cases than an algorithm that does nobody's trying to do it. So, of course, it caused a lot of controversy and discussion, and it got us a fair amount of attention, I think, in the field and it was from that this discussion that happened in the field that it started to gradually dawn on us. I think that the discussion goes beyond the field. People would say things sometimes I've been known spoken about it, people

would say, “well, that’s interesting, like algorithmically, but what does it actually mean for me? Like, I have a lot of objectives, like does this apply to what I’m thinking?” And it kind of becomes obvious that this is a general critique of objectives. It’s not just about like algorithms. And gradually, we start to see that’s actually maybe even more important at some level, because this is a ubiquitous, I think, cultural mistake that’s just like everywhere, and it’s not just that it’s bad for society it’s actually personally individually hurts people. Because you can tell people are suffocating in objectives. Most people are not happy and this is something that I started to learn by talking about this, that people just find it cathartic to hear that they don’t have to have every single thing they do have to be justified through an objective, which is actually stifling and suffocating. And so I think like, it started to be clear that people wanted to hear this message like this is actually important, personally, and culturally. And Joel and I had lots of discussions, I think before the book about, should we try to do this because it’s kind of a crazy thing to do. I mean, we’re like, basically computer scientists, we should we try to write this general book to like the whole world? Or maybe there’s a better way of just talking about it, or what’s the right thing to do? But I think we eventually converged on this idea that this is actually super interesting and just, we cannot just leave this on the table, it’s too important to have this discussion.

Peter: So, did this start out as an experiment in computer science, or human computer interaction, or psychology or something else? What was the genesis of this?

Ken: Well, the original PICbreeder, where the idea starts to appear, it was a computer science experiment. I mean, it was people were involved, though. So, it was about computers with people in the loop – it’s about open-endedness really, about discovery. We’re interested in seeing people discover images and we’re interested in the underlying phenomena and what they might tell us about how the world works. So, I thought that or how the world works as in how algorithmic search works. Like, I thought that we would discover something, I wasn’t sure what though. So, it’s really computer science, we didn’t know I didn’t know that we would eventually see, like, identify some problem with objectives is such a general thing. I had no like ambition to like shatter the objective paradigm or something like that. That was just a surprise discovery. But it evolved from there, because once we saw this, like objectives were actually bad in PICbreeder, we start to extrapolate and maybe I should turn it to Joel, to talk about them like beyond that, how that eventually leads into what our thinking was about, like the social aspect.

Peter: So, the beginning of the PICbreeder program in the context of what were you initially trying to do and what was the goals of that before it evolved into the learning about human philosophy that you seem to have landed on?

Joel: Sure, yeah. So, yeah, when I first arrived to UCF to work with Ken and he kind of pitched the relationship between PICbreeder and the possibility of an algorithm. I think, I first thought this was pretty crazy. I was excited, but also kind of not sure like how this would actually work, because my mind was not in the space of the deceptiveness of objectives and personally, I think I was without a philosophy of search in my own life that would be aligned with this. And so while I was maybe excited, and initially skeptical, kind of the work that we did together to change that into an actual algorithm, began to affect not only my perception of machine learning, but also my

perception of my own life. And through that process of kind of asking the question of the pervasiveness of deception of these objectives that lead you astray, within machine learning, you begin to ask the higher level philosophical question of how this applies to searching in my own life. And that I think over time contributed to our discussions and where we're thinking about and leading to the actual recognition that we're saying something using computer science as a grounded philosophical kind of thought experiment, to say something broader about just search in general.

Peter: That led to this realization about the human condition, and it goes very meta, here, because it seems that the program itself that you were doing demonstrated the philosophy that you ended up extracting from it, which is that it surpassed its initial goals and revealed the way you wanted to go instead was somewhere that you hadn't been planning to go in the first place.

Joel: Yeah, that's right. That the algorithm itself, through searching for novelty in this case, yeah, manifested that principle just kind of in an empirical, scientific way and it's, I think one benefit of this kind of computer science type of approach is that you actually have to put your money where your mouth is and create an algorithm that actually does something. And while some people have maybe said similar things in a more rhetorical way, or qualitative way, about the nature of innovation, you can actually really embody this in a way that's repeatable, testable, and you can modify all the parameters and kind of figure out really what assumptions in this case, like does when does searching for novelty actually lead to better performance or discovering more things?

Peter: So, maybe you could describe the initial experience with PICbreeder where you discovered this principle that trying to produce an outcome was counterproductive and I'm going to ask you to correct my characterizations here because I feel like the phraseology is very important and I'm probably not landing in the right place, when I try and get there because of how subtle this message can be.

Ken: So, what we saw in PICbreeder was kind of mind blowing to me, in demand to help to give a little context of what PICbreeder was like, so you can understand the PICbreeder system was basically an online system, it was like a website, where you could breed pictures and so if that's sounds strange, and you haven't seen something like this before, basically like breeding dogs or breeding horses, except you're breeding images. And so you would see some images, and then you can pick one, and that becomes a parent, and will have children, which are like offspring images, which are just basically a little bit different from the parent the way like, if you had children, they'd look a little like you, but they would also look a little different from you and so you could do that with pictures basically online. And the underlying technology was of interest to us at the time, because it was kind of evolutionary computation, genetic algorithm type of technology, special encodings, that were supposed to work in certain ways that are reminiscent of DNA; stuff that the user doesn't need to understand, but that to us was interesting. But it allows you to do this picture breeding and so you could actually under the surface, there's a kind of like an artificial DNA behind these pictures. But what was really interesting was once we

allowed people to go off and do this on the internet, so basically, it was like, we were crowdsourcing the Internet to breed pictures in this giant space of possibilities, we started to see lots of really cool discoveries, like people discovered things - because as you can imagine, they start as total blobs, doesn't look like anything - and eventually, we will discover things like butterflies, and cars and skulls and planets and it just like really remarkable that through what we understand is random mutation. People over very small numbers of iterations or generations, we're discovering things that look like real things reminiscent of the real world. And that on its own, I think is interesting, but what was really interesting was when we looked and saw how they were doing it, and the reason we could do that is because we have an entire history of everything that everybody ever did recorded. Which I think is fascinating, because there's no artifact like that, that had previously existed that I know of, where there was like this giant phylogeny of search, you can evolutionary tree thousands and thousands of images; items where the entire history is actually recorded and known. So, you don't have to, like infer what happened and so looking at that history, what we noticed was that people who found things like the butterfly, or like the car - actually I was the one who found the car. So, actually, interestingly, the first realization of this was when I found the car, because I actually experienced for myself this very strange phenomenon. And what we found was that it's always the case, at least 99% of the time, that the intermediate stepping stones, the images that people went through to get to what they found, don't actually look like the thing that they found. And this is really, really fundamental. Because if you think about it, what this means is that those that those people would not have found those things if they had been looking for them, because they wouldn't have chosen those intermediate stepping stones, because they don't look like the thing that they ultimately got. Like, the only reason they chose the intermediate stepping stones is because they *weren't* trying to find the things they ultimately found. Because they were interesting in their own independent right.

Peter: Right, because they saw an image; they thought that looks interesting and it took them down a path to something that was fulfilling toward a goal of some kind, satisfying, but didn't look like the thing that they initially said was interesting. And if they had been trying to get the thing that was the end result, they would have rejected that first interesting thing, because it wouldn't have looked like it was on his path to get there.

Ken: Exactly. And one just small point to notice that sometimes people were branching from each other's discoveries. So sometimes it was like somebody discovered this one thing, and then someone else say, "oh, I'm going to keep breeding that." So, sometimes there's multiple people involved in this chain. But that was the phenomenon we observed and to me, it was totally mind blowing, because it just contradicted everything about what I had learned about achievement, and engineering, because things are always like taught to you like, you set your objectives and then you basically make progress towards that objective. Maybe you do top-down design, you do bottom-up design, there's like different methods, but what you're trying to do and that's what we're able to do and this is like the exact opposite. It's like, the only way you can do anything is by not trying to do it and I think it maybe is like, it doesn't become maybe the reason it hit me so hard is because I experienced it. If you just hear that, or if you just see it like as a phenomenon, like shown to you as a sequence or something you might be like, "ok that's interesting", but

when I actually went through it myself, because I found this car, and I knew I wasn't looking for a car, I actually experienced this. I choose an alien face and it turned into a car because the eyes turn into wheels. It made me feel it it viscerally like how this is totally bizarre, because it just struck me that I could not have achieved that, like found this car, if I had wanted to do that. The only way it could have happened is if I had been not looking for it. And just the feeling of that was like what this doesn't work the way the world's supposed to work, or the *world* doesn't work the way it's supposed to work and then the big question after that was like, Is this just about PICbreeder? Or is this like a bigger point? And there's it turns out it is a bigger point like it does. It's not like there's something special about PICbreeder that makes it only true in PICbreeder. This is the way all complex spaces work. Because for something to be a hard problem means you don't know what the stepping stones are. If the stepping stones were obvious, it's like a tautology, it's not hard, you know, because like then you would just solve it. So, like it has to be the case that the things that lead to the things we want to achieve, like say cure cancer, create AGI like whatever, some super ambitious things are things that are not like those things. If they weren't, we would just do them. And so this is just the general structure of search spaces that are complex or ambitious and so that realization was like, wow, this is profound in terms of its implications and that's when Joel just suddenly appeared, because Joel just started right as I was having this revelation.

Peter: Joel, what was it about this that pulled you in at this point?

Joel: Well, it was, I think, the provocativeness of it, and the breadth of it, those things kind of spoke to me. Honestly, when I was starting my PhD, I think I was a bit naive, I didn't really know what I was getting into in terms of like a graduate kind of program. Like it was an idea maybe what that would be like, and so there's some serendipity even in that process of coming to this lab with this person who's telling me all these kinds of really ambitious sorts of ideas. And so I think that yeah, for me, it was a bit kind of a learning process of getting to some place where I actually could begin to implement these ideas and get a taste for myself, what actually was going on. And I do remember one kind of impactful experience for me it was when we had kind of decided on the first kinds of the experiments we were going to do with novelty search, like a robot in a maze and there was like an all-nighter sort of coding session where I got the first version of the algorithm going and as the sun's coming up, the first results are kind of coming in and yeah, that the deception was being shown and that it was kind of working in a way that we thought. So, kind of like that was maybe a similar moment for me to as to when Ken had his alien face into car moments, I think was just seeing the first results come in and like okay, well this, this actually works, which is pretty incredible.

Peter: So let's try and put some boundaries of parameters around this for the help of the audience who might be thinking that this is counterintuitive; if my goal is to go to the supermarket and get a loaf of bread, I'm not going to approach that by going to the park and throwing a Frisbee or picking some other random action. That is something I know how to do and I know how to optimize that that task. But in something that we don't know how to do like cure cancer, now it seems very much productive to explore directions that make no sense. Can you help put some guidelines, some rails around that?

Joel: Yeah. I think one principle you can use to think about that it's just the notion of stepping stones, that for many simple problems, like going to the grocery store, the stepping stones from here to there are clear, you've maybe done it before, it's pretty predictable. And then when it comes to problems, of some ambition, when stepping stones to get there are not obvious. You can say more than one stepping stone away. So, things like inventing something that has never been invented before. There it's quite unclear how to get there. And I think the tricky thing is that I think our intuitions about stepping stones are often not great. And that's part of the problem, that we have maybe sometimes we're overconfident that we can get to someplace and that some things that maybe seem like they should be within the domain of just pushing towards an objective and achieving, may actually not be that way. So, things even in our own personal lives about finding a good relationship or getting a good job, sometimes those things, they're at the edge of what we might think of as the realm of the ordinary objective-driven stuff, or that will work and the realm of where actually it's quite ambitious and exploration is really needed.

Peter: So this seems to that you've opened the door here, into a realm that few computer scientists have stepped into, which is management/leadership guru presenting this provocative novel, counterintuitive idea. Did you step through the door? Did you become that kind of messenger, then, to take that to the people that can use that without it being a lesson in computer science?

Ken: I think we did. I mean, that was the sort of the role of the book was to start to take this to a broader audience and it was, I mean, even before the book, I mean, one of the reasons that like it was start to be clear that this would be effective to do that was because I was starting to talk to audiences already that were not computer scientists and noticing that this resonates a lot maybe more than computer science sometimes. You know, like I talked to the Rhode Island School of designers. It's really early experience I had like, because someone had seen me at a computer science conference, who was a student at the Rhode Island School of Design, which is basically like an art school, art and design and so he said, "I think the people at the school would really be interested in what you're saying." That was a really interesting kind of test case. Because I had before that really spoken to formally two groups that are like way outside computer science. It sounded really interesting. I could totally see why it might be interesting to know, but I didn't know what to expect and it turned out that the artists were just extremely moved by the whole message, because for them they were in a world where a lot of their lives was defined by people not understanding what their objectives were, like their parents, like, why are you doing this is pointless get a real job do something, it's like a real career, or even their teachers are like, well, that's nice, but what's the point, what is the objective? What are you trying to do? And the reaction to this message was so emotional and cathartic that it just like wow, was this is I could just tell them that this is going to be important, like all over the place, like it's not just a computer science topic at all. I had, they gave me private meetings, where the students could sign up later to talk to me privately. And those meetings really affected me because I was, they were almost like therapy sessions. I didn't know what to expect when people come in and gave their whole life story and they're practically crying in some cases. And I was just like, this is so incredible. Like, this is computer science, like, this is boring dry stuff, like algorithmic stuff that Joel and I were working on it these people almost coming to tears. I mean, when has this ever

happened? That was when I was like this definitely has to be a book like this is so interesting. And so after the book, like more and more of these kinds of opportunities, start to pop up when people read the book, or, or some presentations about the book and they would be like, “yeah, we want to hear more about this, get your opinions about this” and so by now, seven years after the book came out, like, I’ve been able to be exposed to a huge variety of communities it’s really interesting and learn how they think about it, that there’s and what this means to them.

Peter: What does that arc been like for you, Joel? What places has it taken you that that are new and unexpected?

Joel: Well, I think the book was a start to really dive into the more philosophical nature of search and of computer science, I think that’s a thread that I think has really stayed alive for me. So, just that computer science can be a vehicle to ground out things that are otherwise unclear. So I think the other message of the book being about how search relates to innovation, and how objectives can be misleading, is really intriguing and I think by going beyond that it’s just interesting, in general, that almost any kind of qualitative phenomenon, in theory can be boiled down into some kind of thought experiment within computer science and so some of the directions I’ve looked at beyond that are kind of oddly in like the moral philosophy space trying to bring together computer science and moral philosophy. And now my interest is more kind of in your science and psychology, and looking in that direction. So, I think that this was a really profound experience for me to be involved in exploring how rich computer science can be, in terms of like actually exploring questions that are relevant to people’s lives. I think that’s the kind of thing that has really shifted my trajectory.

Peter: It seems the ultimate validation of pure research; like you could keynote at any conference on pure research and tell the audience why what they’re doing matters no matter what anyone else thinks about it.

Joel: Yeah, I think one interesting thing, there is something that I feel is unresolved in the book that is kind of an open question for me is to what extreme can we take that and is there a balance to that? You know, like, in the book we talked about, as an example of someone who’s just working with Legos, that’s kind of their life’s passion and we don’t know that will lead which is true. But also, it does seem like we do care ultimately about allocation of resources, and how to how to do that and the message the book is not like anything goes in terms of like, everything goes, everything should be equally promising. But to me, there’s still something confusing about, like, if you have a fixed number of resources, like how actually should we allocate that to get to where we want to go, if that’s even possible.

Ken, was there pushback along those lines? Did you get reactions from like business people saying, look, we’ve got budgets, what are we supposed to do with this?

Ken: Yeah, I did get to kind of learn how business people or people in positions of power will actually sort of mentally process these messages and it’s not really that way. I like that’s how I feel at least the way I found it’s not so much that people are like, no, we have to do these things like we don’t, we can’t do that it’s more like, I found that most people who are in gatekeeping

positions really like this, they would be almost always across the board like, this is extremely interesting like motivational, and we should change our processes or something like that. But then what would happen is that we would do, let's sit down and talk about like, they want to get into more detail, like, what can I do? And when we sit down and talk about it, it's kind of like, "well, I would love to do all this stuff, but I answer to all these people, and these people, they're not going to necessarily allow me to do that and I'm worried about this and they don't really understand all this stuff and I don't think they will." So at the end of the day, it's like, well, there's really not much that I can do, I've tracked and so I think it's a little bit frustrating, you know, that hierarchy sort of makes it very difficult to push things out of this static configuration that they're in and get us into a different way of doing things. And it's sort of illustrates why I think you need to have a discussion like this, or why there should be a book is because without this being an object of conversation, those kinds of issues are just going to be stuck. Because it's true that you can't just go to your boss and just say, "hey, maybe we should have less objectives" there's no context for that. So, you're on your own with this radical statement, and it's just kind of really hard to defend. And so it's somehow, I think, a broader social discussion is necessary to dislodge some of this. Or else, it's just going to be impossible to break out of it.

That's the end of the first half of the interview, which as almost always we split into two so you have more digestible chunks for each episode. Believe it or not, that's the first time Ken and Joel have been interviewed together, despite delivering some pretty high-profile keynotes and the like.

If you want to see Picbreeder in action so you have a better idea of the program that started all this, if you google for 'picbreeder' all one word, it should be the first hit, at nbenko1.github.io, and you can try it out for yourself. Just select the most interesting image in each cycle and click 'Mutate' and see where it takes you.

In today's news ripped from the headlines about AI, in the UK, Channel 4's alternative Christmas message was generated and delivered by the Ameca robot, which we mentioned a few episodes ago. Ameca, from the company Engineered Arts, has the most lifelike facial expressions – they're definitely flirting with one side or the other of the uncanny valley – of any robot I've yet seen, and its GPT-3 back end enables it to act as a customer interaction agent at Dubai's Museum of the Future. The content of the speech was generated by GPT-3 and no surprise that it was very convincing, this is exactly the sort of content that is easy to generate by a transformer now and not taxing. So it's now not even remarkable that it – or she – came up with the text "Be 'neither happy nor sad' about the past year and 'take it as a learning opportunity, a chance to change the way we think about the world and a reminder to help those in need whenever we can.'" I'm sure we'll be seeing and hearing a lot more of Ameca.

Next week, we'll conclude the interview with Joel Lehman and Ken Stanley, when we'll talk about how to apply the principles they learned in practice in your business and your life. That's next week, on *AI and You*.

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

<http://aiandyou.net>