



# OPEN AI PODCAST PT 2

## AUDIO TRANSCRIPT

**Fernando Lucini** [00:00:00] So everybody, welcome back to our conversation with Peter Welinder. Let's get right into it with our first question. We need to talk about humans and the dissection of humans with all of this. To some degree, we're talking about this, and we love this stuff with scientists and engineers. We love what it's trying to do, and we're openly saying that we're letting the world figure out, you know, what it all means and it's a critical part of the evolution, right? But at the same time, we are both OpenAI and myself in Accenture, we are quite, you know, conscious of the role that we as humans play in all of this. We've talked about a bunch of examples there, there's examples in education, there's examples in understanding, by the way, the example I tend to use the most is if you think about most companies today, they're very knowledge heavy and that knowledge is generally represented in digital text and funnily enough, generally attached to human beings and individuals and not really goes any further than that just because of the limitations of comprehension of that. So does opening up... Let's say we can use GPT-3 to understand an entire gas and oil companies' blueprints and all this stuff.

**Peter Welinder** [00:01:24] Yeah.

**Fernando Lucini** [00:01:25] What does that mean to the human beings? And my argument is - and I'm going to let you do your part on OpenAI - but my view has always been that, like everything, it's a multiplier of our mind. That's the point of it, right? It's trying to multiply our mind. It doesn't change because no matter how many times Fernando Lucini, electronic engineer and data scientist and software engineer, no GPT-3 in the world is going to make me a petroleum engineer. It is what it is. That's not what it's going to do. But will it help me access the information more relevant and quicker so I can actually do whatever it is that I need to do that requires me to learn a bit more about it? So to some degree, it's worth pointing out, how do you guys come at it from the perspective of what is the role of the human in creating something that moves

towards general intelligence? Where's the human in OpenAI?

**Peter Welinder** [00:02:18] Yeah. No, that's such a great question. And you know, I think we should all recognize that if we get to AGI eventually that will change a lot of things. And I think that's, you know, a big part of the premise of OpenAI as a company is that it's important to set it up such that the economic returns eventually goes to all of humanity because it will be disruptive. But I think, in the short term, what we will see is, to your point, it's really kind of really aiding and augmenting humans to do just more of the work that humans are really good at. Like, I like to think about. There are three things that these models do. The first thing is that now they can aid your creativity. The second thing is that they take over more of the mundane work, like the work that most people are not so excited to do anyway. And lastly, it kind of improves the work product. It kind of actually makes the output of what you're doing better. So if you look at the first one like creativity, one of the first applications of GPT-3 in real products that people built on top of API was in authoring tools. And it really was a way to kind of help writers with writer's block, you know? If you want to start, you have a character and you want to describe what this character looks like or something like that. you have to be creative. You have to come up with like, what are the ways this character might look like and there was a number of tools built where you could then have GPT-3 generate like five possible ways the character could look like and then you sample those and that kind of unlocked you. You could take one of those and you could edit it, or you could like sample more, but it's got to unlock to you. Like we had one author who told us like, hey I've been struggling with writer's block for kind of half a year now I have written like ten pages. Then I got access to GPT-3, and I finished my book in two weeks. It's amazing.



**Fernando Lucini** [00:33:24] Oh wow! GPT-3 is not writing the book, right? This is not what it's doing. It's helping with imagination. It's helping. It's not replacing them at all.

**Peter Welinder** [00:04:31] It's a brainstorming partner. Like it's essentially a colleague or a friend that you're brainstorming with. And so it really helps you kind of just supercharge your creativity. Obviously then it can do a lot of the mundane stuff, like for example, for Codex, for our programming models, you can actually, through this project called GitHub Copilot, autocompletes your code. A lot of this sort of code helps you with, if you like writing more boilerplate code, for example. Or you want to like document your code. Like a lot of people, a lot of programmers are not very good at documenting the code because it's like, I just wrote it. Now I have to explain what it does. That's a bit boring. But then you can just like have the model do it. The model can look at the code and write the documentation for it, right? So it takes over needs.

**Fernando Lucini** [00:05:20] Spectacular. Let's stop there for a second. I mean, that could be a generational change for us software engineers, right? We love writing our notes, right? Oh, that's fabulous.

**Peter Welinder** [00:05:29] So it can help with those things that are like a bit mundane but like can add so much value. And then I think the last thing is around making you better and I think this is where, similar to codex, people are using it to actually spot bugs in code to kind of refactor it to be easier to read. It's similar for text. It can help you remove all of the cliches from this piece of text that you have written or make it more engaging and so on. And the model can transform a piece of text into a different piece of text based on the instructions you're giving it. So it's really helping you create much better versions of what you've already created.

**Fernando Lucini** [00:06:23] So let's go back to the application of the mind, right? None of these examples are doing anything except amplifying, bettering. I'll give you one that I like. I don't know if happens to you, but I search Wikipedia 20 times a day.

**Peter Welinder** [00:06:37] Yep.

**Fernando Lucini** [00:06:38] If you're endlessly curious about the world, every single one of my interactions with a browser starts with wiki space and whatever.

**Peter Welinder** [00:06:47] Right.

**Fernando Lucini** [00:06:47] Endlessly. Endlessly. And I keep on using the simple example of well if you turn that off for me now, I'll be in a state. I wouldn't know where to go because it fills a special place in me where knowledge is really important.

**Peter Welinder** [00:07:02] Yeah.

**Fernando Lucini** [00:07:02] So if we then multiply by this and you say, well, is GPT-3 and it's use cases the next generation of this fulfillment of knowledge and fulfillment of things with action rather than answer, right? As you say, the great example of, I've got this paragraph. It's not very engaging. What a difficult thing to say in terms of, you know, how do you even quantify that? Can you make it more engaging? Or the other example you gave, I'm going to use that one as well. You know, remove the cliches. What a wonderful what a wonderful one, right?

**Peter Welinder** [00:07:31] This is such a good example because actually we've had people tweet about how they started using GPT-3 instead of Google to find answers to things. And it's not only because they just wanted a straight answer that might be hard to get from Google, but it's because once they get that answer, they can have a dialogue back and forth with GPT-3 to kind of go deeper on the topic. For example, some of the coolest applications I've seen is if you want to have a conversation with... As engineers, what's a famous person that we all love? Richard Feynman. If you want to have a conversation with Richard Feynman, you can basically set-up GPT-3 to help you with physics, explain some physical phenomena as Richard Feynman would do it. You know, that's just so incredibly powerful, right? Like you can kind of touch these subjects much deeper than just from reading about them.

**Fernando Lucini** [00:08:27] I'll tell you what. I write the worst tweets in the world. It's not my thing. I prefer the reading than I prefer the outgoing. So I'm going to ask it to make my tweets more interesting.

**Peter Welinder** [00:08:37] There you go.

**Fernando Lucini** [00:08:38] I'll tell you what. It'll make my LinkedIn posts more interesting. I mean, I know we laugh but the truth is this is material stuff. And we should talk about DALL-E for a second because it's really breaking the mold even further. So the next thing you've got is - and I'm fascinated by this and of course, I work with you guys so I'm a bit maybe more educated than the average - but the ability to create images, versions of images from narrative. I mean who woke up one morning at OpenAI and thought, we should be doing that. We should be using natural language, and it should create images. That sounds easy. So where did that come from? What was the nexus of that?

**Peter Welinder** [00:09:27] Yeah, you know what we really want to build at OpenAI, as I said, is artificial general intelligence. And so it's ultimately models that can understand the world much deeper like humans understand it. And I think what that really means all around are multi-modal models, models that can deal with text, with images, video, audio and all those things together. And, if you have that, that just provides a much richer understanding of the world and also kind of enables, I think, interfaces that are much more natural for humans. Sometimes it's going to be much easier to explain a concept to another human if you can draw something, right? Like how many conversations have you had where you're in front of a whiteboard and you draw something, right? That's kind of really where we're going towards here. But we're also scientists, so the way to kind of get started is you start in some way. If you actually want to go to models who can do all kinds of things, then you want to start with models that can do text and then you start, okay, like how can we now go from text to images? We actually started in the other direction where we went from images to text. We had these clip models that basically we embedded both text and images, and you find that you can compare them, their similarities and so on and build classifiers on top. But, but yeah, it's really towards that just richer representation of the world, deeper understanding of the world, really see the limits of what these models can do.

**Fernando Lucini** [00:11:07] So I was trying to explain this to my ten-year-old. So hold onto your hat, Peter, this is how I explained it. I was telling him imagine you're in a place where you can smell the swimming pool because you can smell the chlorine. You can see the swimming pool, and you can do it with incredible precise - that a swimming pool, not a river. And you can hear it. And, of course, you can extend and say touch it, but let's just do those three. I explained to him that we can get the machine to do one of those three things, but we can't get the machine to do all those three the way that you're doing it, which is all of the things come together in this splash of things.

**Peter Welinder** [00:11:41] That's right.

**Fernando Lucini** [00:11:42] I won't tell you what he said to me because it was insulting, and it doesn't speak well to my children. [Laughs]. But so not quite the objective is everything goes towards general intelligence, but it's in the context of trying to think about general intelligence. It's a step towards actually multi-modalism, which is such a complicated thing. Let's talk about some of the use cases. We're talking to some of your folks who help me all the time and bless them. We were talking about marketing and marketeers and creators, and we had another one of the examples of humans and how humans fit into this because it's, hey, does it replace an artist? And the idea is no, absolutely not. Artists are artists. Art is art. Photography is photography. All these things are perfectly natural things that happen all the time. But does it support things like speed or, the example we were using is, well, I've got the image, but I need it to be slightly different in these five different ways. For example, do a different background on it, put a different wallpaper, or put a chair somewhere or, even as you said, when we don't have artist at hand to be able to create things in a very quick way. What are the interesting things that you've seen in that area that are worth mentioning?

**Peter Welinder** [00:13:08] Yeah. And one thing I would caveat is that this is still super early. We're really starting, only a month ago, so we're still learning a lot. But I think ultimately the way I look at it is that is going to really change the way I think a lot



of artists work. And it's also going to enable a bunch of applications that just weren't possible before. So we talked about those in turn. I'm really doubtful that this would actually end up hurting artists. I think it would provide yet another tool as an artist. And in fact, the artists that we see using this DALL-E today are extremely excited about it because it enables them to actually explore many more ideas, much faster, where they would basically before limited by the tools they had. And so it just gives them a much bigger creative outlet where they can then take that output and kind of go towards kind of the vision that they have in their mind. And it also allows them to kind of create new sorts of art that just they couldn't create before because of the scale of which they're able to create say individual images and so on and transform them. It's much faster, I think. But I think the second thing is that, to your point, say in marketing and a lot of these areas, there's a bunch of cases where I think if you wanted to create something customer, some custom artwork, some custom kind of photo or something like that to market a product or market something else, then you didn't have many options. Like most people can't hire an artist. Like if you're a mom-and-pop store somewhere like you're not going to be hiring an artist probably. You probably wouldn't even know where to start if you wanted to do that. But if you have something like DALL-E, then you can actually get something good enough pretty, pretty quickly. I think similarly, we'll see probably tons of applications of this in entertainment where you can now customize stories, the assets around stories, like the imagery. Eventually, probably, the way this field is moving, I wouldn't be surprised if we have like videos and so on generated in just a few years and so it creates much more personalized and custom experiences. I think we'll see a lot of entertainment and games and so on. And I think in fact there are already people that have created like music videos with still images and so on, but they put together like basic music videos with these things, people that are creative but not really good at drawing can suddenly kind of create something really beautiful that they just weren't able to create before. And I think this is one of the critical things around these new models that GPT-3... it can do certain things better than me. I'm terrible at writing poetry, but I would say like, you know, GPT-3, especially the business models, are quite good at poetry. But for like creating images,

art, drawings and so on, DALL-E is so much better than me. Like I just cannot do this stuff, and suddenly it enables probably 100 times as many people to create things in a space that you just weren't able to create something before. And I think that in itself is just going to kind of unlock so much possibility in terms of what people make with it.

**Peter Welinder** [00:16:39] Go back to the multipliers of the mind, right? Which I think it works for me perfectly. Artists are great examples of every tool is a tool. Let's use all the tools to create. Creation is great. Let's do and create. You're absolutely right. We haven't talked about education, but it's an enormous, enormous potential of things like DALL-E in education, where we're constantly representing, teaching, showing, and you break down some of the barriers to be able to connect the things we see and describe, which are complex things, versus the things we see. So I was doing some fun things with Jacob, with my autistic boy. This is the best thing in the world, the best thing he's ever seen. What was the example? It was the whole backpack that looks like an avocado thing. We laughed for so long. So long. The amount of combinations of things he could think that would be the best things in the world, which are his form of art, were absolutely fabulous. And of course in a more serious business world, it happens all the time. We're visual beings by definition.

**Peter Welinder** [00:17:49] Yes.

**Fernando Lucini** [00:17:50] And what is next? And by the way, we should make the point again very clearly, this thing is a month old.

**Peter Welinder** [00:17:56] Right

**Peter Welinder** [00:17:57] It's a month old and GPT-3 is two years old. We're talking about it like it's been around forever. It is the beginning. But that said, because I think it's a good way to tie up the conversation, given all this, what's next do you think? What can we do? And again, when we talk about all these amazing things that we should be cautious, but what do you think is next for you guys?

**Peter Welinder** [00:18:24] Yeah, it's a good question. I think this is an incredibly exciting field because it's just moving so fast. These kinds of generative models



were able to generate small gray scale kind of faces five years ago, and today they're generating like photorealistic images of almost anything. So the pace of progress here is kind of incredible. And I think really for us, one of the biggest focus is just kind of making these models better and easier to use because, we think there's just incredible potential and we want a lot of things being built on top of it. Like, for example, I think education is one area where I'm extremely excited about people building things. I don't know if you heard about this concept of the Bloom 2 sigma problem. It was a study back in the eighties from a guy called Benjamin Bloom where he showed that if you have a personal tutor for a student, they will learn at the pace that's two standard deviations higher than a group of 30 people with a single teacher. And I think, what AI can do is essentially provide a personal tutor for everyone. It's just like somebody needs to build this. And I think the state at which this technology is today and where it would be in two years, like with this pace we're at, it should be fully possible to kind of start building these sorts of solutions. But, going back to one of the first questions you asked me in this chat is why do we do an API? Well, it's very much because there's so many applications for this, and we don't have the ability or expertise to build all these applications. So our focus is really to kind of just get these core models better and better, smarter, and smarter and easier and easier to use. And keep on bringing down the prices for access and so on to make them just as accessible as possible. But I think, you know what you'll see more of is models that can do more things at once. We have now models that can do code and we have models that can do text, but our latest models are quite good at both code and text, and eventually there will be good at images and so on. So they will just be able to be smarter around a bunch of different kind of capabilities, and then we'll just see what people make with them.

**Peter Welinder** [00:21:12] Absolutely. And I'll tell you what, it's a great ending to say that. I'm a father of three and through COVID, I would have needed your tutor to supplement. And we should say there are wonderful, wonderful teachers and tutors and people that teach our children and teach society as a whole. But even then, we all know that we could use something like what you just described. So I'm totally like you, anything that supports our growth in

knowledge of our tutors and our learners is amazing. So with that, I'm going to thank you, Peter, very much for coming and geeking out with me for a while on this wonderful topic. It's been great fun, and I hope everybody takes away some of the examples Peter has given us. I think it helps educate all of us into how some of this stuff can be used in a creative and interesting way. And I actually would tell everybody, please go, and use OpenAI and join in the journey, right Peter?

**Peter Welinder** [00:22:10] Thank you so much for having me. This was super fun. I really enjoyed it.

**Fernando Lucini** [00:22:15] My pleasure. So thank you very much for listening to us. Please subscribe. Please come and see us and we will talk to all of you soon. Thank you very much.

Copyright © 2022 Accenture  
All rights reserved.

Accenture, its logo, and High  
Performance Delivered are  
trademarks of Accenture.