



# OPEN AI PODCAST PART. 1

## AUDIO TRANSCRIPT

**Peter Welinder [00:00:00]** Makes kind of really advanced AI accessible to way, way more people. That's also what we've seen in terms of our user base.

**Fernando Lucini [00:00:15]** So hello, everybody. I'm Fernando Lucini, a chief data scientist at Accenture, and today I have the distinct pleasure as a massive MLP geek to have Peter Welinder from OpenAI. We're going to let him talk to you a little bit about OpenAI and its hopefully an excuse to geek out for quite a long period about how the world has changed and what OpenAI are going to do to further change. And were not going to stick to just text, are we Peter, we're going to go beyond that. Thank you for being here Peter.

**Peter Welinder [00:00:44]** Thank you so much for having me. Super excited to be here.

**Fernando Lucini [00:00:46]** No, as I said, it's exciting for me and for those who know me know text and trying to get interesting things out of text has been my life for a very, very long time. So suddenly in the last year since you guys started, suddenly I started thinking, oh my God, the world has truly changed. All the things we used to do now this is solving some of those things that we thought we were solving. And so it's actually super exciting. And I know that everybody listening will be super excited about the journey and stuff. I'll tell you what, it's not a bad place to start, right Peter, telling us a little bit about yourself, which obviously everybody's always super, super interested in and to tell us a bit for those that don't know about OpenAI as a company, its beginnings, maybe bring us to the present, I think maybe. Is that fair? Is that okay with you?

**Peter Welinder [00:01:32]** Yeah, sure. I'm super happy to talk a little and give a little bit of background on OpenAI, kind of how it came about. My official title is VP of product and partnerships with OpenAI. What that means is really that I'm in the really fun position of taking the amazing output of our research team and bringing it to the world as products, primarily our API. And I'll talk a little bit about that. OpenAI is sort of ...we call it a research and deployment company. It was really founded to be a research company with a mission to build AGI - that's artificial general intelligence. And what that means is really general-purpose AI can reach and eventually surpass human level capabilities on most economically useful tasks. So, you know, you can think of AGI as sort of systems that can do the work of highly educated humans, like lawyers, engineers, doctors, scientists and so on. But then the other part of our mission is, you know, you can imagine if you build AGI, that's going to be a very powerful and really world changing technology. And so OpenAI was founded with the mission of doing that but making sure it's built in a safe and an ethical, responsible way and then deploy it in a way where it can benefit all of humanity. And so that's the first part of what we call our company, the research company. The second part is deployment company. And that's primarily what my focus for the past few years is on the deployment of this technology into the world. So, the way we do that today is we have an AGI platform which lets people access this technology that that we are building so that the first thing that we built that was really useful to be deployed through an API was this language model called GPT-3. It is essentially a language model that you can give instructions and you can kind of get text back. We eventually developed



Codex, which is a similar language model but focused more on code. And most lately we also developed DALL-E, which is a model that goes from text to images. And so that's a big part of what we're doing at OpenAI today. In addition to the core research to making our models better and better towards AGI is to make sure that it's available to be used by as many developers as possible, because we think there's huge amounts of value, a lot of new applications that can be built on top of this technology. And we want to make it as easy as possible for others to build on it.

**Fernando Lucini [00:04:20]** It's interesting because as a product - researcher turned to product guy, GPT-3, must be the worst product name ever as products go, right? Whereas then DALL-E's a fabulous name which makes all of us smile and for the listeners which make it clear GPT-3 with its fun name, certainly its objective is to understand the written language in a way that we haven't seen before using machine learning or AI by using an enormous, enormous amount of information and then very clever science to give us that comprehension and a bunch of other aspects like comprehension, summarization and a bunch of other things surpassing almost like a step change from what we had before. Right. and the way I tend to think of codecs is, well, if you've managed to do that with language as complex as human language is a programming language is a language.

**Peter Welinder [00:05:13]** Right.

**Fernando Lucini [00:05:13]** So that sounds like a sensible thing to do. I will talk more about them. DALL-E then it becomes for me, another fundamental jump, which is, well, you know, can we tell a machine through descriptive language what we want, and could it create an image of what we have asked it. I think that requires another level of geekery that I think we will get to. But it's an interesting set of jumps, right? A couple of questions for you, because, of course, you create this this incredible science. You guys create some of the most fabulous reinforcement learning out there. And, of course, we will clarify

some of these, but you do it in the shape of an API.

**Peter Welinder [00:05:56]** That's right.

Fernando Lucini [00:05:58] Which is a choice, right? Talk us through that choice. Is it a choice? Because these things are so big and actually it doesn't make sense to actually try put in a way that you package it? Is it just a safer way to deliver it so everybody gets the most benefit? Why? And I ask you this totally, openly, open agenda, Pete, because of course, our own customers and people out there ask the same question. Is this something? How do I use that inside my firewall? And I think it's worth hearing the answer.

**Peter Welinder [00:06:25]** Yeah, for sure. And, you know, this is a great question, because when we first built GPT-3, you know, before that, as you alluded to, we had done a number of other projects at OpenAI towards kind of really showing how we can push by scaling models and using tons of data, you could get really impressive results and we had shown this across a few different domains. For example, we beat the world players on the kind of a well-known computer game called Dota 2. We had a robot hand that could solve Rubik's Cube and so on, but those were kind of cool demos, but they weren't really at the level where they were useful. And when had developed GPT-3, that was the first time we had something where we felt like this this seems like it could actually be really useful to build a real product around. And that's sort of why we, with GPT-3 started our kind of the more deployment path of our company. Now, to your question, you know, we actually debated a number of different ways we could deploy this in the world. We were thinking about execute build like a really good translation system, some kind of, you know, more virtual assistant like experience. You know, as you said, it's really good at things like summarization. And so like we were kind of going through this, different ways we could turn this into a product. And where we landed was basically it was really hard for us to make up our minds because one



of the really kind of core things about GPT-3 was just how flexible it was. Like it used to be the case where for each of these use cases, like for translation, summarization and so on, you would build a very custom model and you would train it specifically for that use case. And what made GPT-3 different is that you could ask it or to kind of do any of these tasks, and we would do it at like state of the art or close to state of the art. So it would do it really, really well. And the way you would do it is, and we can talk about that a bit later, but you would basically do this thing called prompt decide where you would show it a few examples of what you wanted it to do than it did it really well. And when we kind of, you know, when we really thought about the generality of this this model, we thought like, well, you know, there's just so many different things you can apply this to. And, you know, we would be able to think about a few problems that kind of make sense for us and might have big, big business value and so on, but there's just so many more problems that other people know about that we don't know about. So actually the way to kind of make this maximally useful for everyone is to make it available as an API so that anybody could kind of take this technology and see like, what problems do I know about that I want to solve? And, and they could just apply it to the problem. And that's what we've actually kind of seen in practice is like there's tons of applications of GPT-3. which I don't think we had really thought about when we did that initial kind of brainstorming about how we could apply it. So I think in the end that's been a really, really great choice, I think.

**Fernando Lucini [00:09:30]** Yeah. And I was going to say is exactly that. If you maybe, for the listener, put that into perspective so you, you've created this, this, this piece of science and technology that has effectively a great comprehension of the world that we're living through text. Then the next question is, what do we do with it? And I think you're totally right. It's such an open question and there's so many things you could do that immediately think ok well, the places where we have conversations in chapel, it's just an easy one. You think, okay,

that sounds good. And by the way, it's not where we're most fulfilled. It's just a good place to fulfill people where conversations are generally filled. You got the amount of documentation we read. By the way, I'm still waiting for somebody to crack how to use GPT-3 to help me deal with my email. So when you have 5000 emails a day, how do you get something that just helps you deal? You don't have to answer for you though we may get there, right? So there's problems that are really intractable but opening it to the world means that to some degree, you know, you're really opening up to the imagination and the problem setting up everybody rather than set the rules beforehand, and that that's how you use it. This is good for summarization. This is good for classification. It's good for comprehension of language. How would you use it? Right. I think that's, I guess, where you guys were coming from that.

**Peter Welinder [00:10:49]** Exactly. And, you know, and to your point that this is what we see today, like we see people that are, for example, tackling email with these language models. And I think, you know, one thing that GP-3 specifically and I think like one thing that has been really powerful to see is that the people that are doing that, they are developers, but they're not necessarily machine learning experts. And, you know, that's another very interesting thing that we found with GPT-3 was us, you know, it's almost like a new way of programing and doing machine learning where in the past the way to kind of get really good results was, you know, you would collect a bunch of data and then you would train the model on that data, you would get the model out or you would take like an existing modeling and you would kind of fine tune the model, would customize it based on kind of a new data set that you've got. And with GPT-3, the kind of new paradigm was this idea of either kind of zero shot or a few shots learning where you would either, for zero shot, you would essentially just ask the model what you want to do. So imagine like you wanted it to summarize a piece of text, what do you do? Well, you basically, you take the article, you paste it, and at the end of it you write, you know, summarize



the above article and you just give that text to GPT-3 and what you get out is a summary of the article. So like it's text in, text out. Similarly, if you want to do translation, you just ask, you know, translate the following sentence into French and you would write that in English or whatever other language, and then what you would get out is the translated sentence. And oftentimes if that wasn't enough, you would provide a few kinds of examples in your prompt. That's the thing that you sent to the model, and you'll get something out. And the powerful thing here was this is just the fact that to get really good results, you had to be really good at crafting the prompt. We call this kind of prompt design and in fact, you know, before we had done the official public release of the API, we had a few hundred people that we had given early access and one of the best people to get the most impressive stuff out of this was not a like machine learning specialist. It was a guy called Andrew Mayne. He's actually an employee of OpenAI today, but we call him the GPT-3 whisperer. He was just really good at crafting these prompts, and he's also kind of a very accomplished writer. And the way he kind of got the models to do really amazing things is by kind of like writing a story around what it should do, and that was enough to kind of put it into the writing, right? Like mindset, you know, anthropomorphizing it a bit, but like setting it in the right way so that he could really output what you wanted. So what we saw is that the people that managed to get the most amazing things out of the model were not the people who were machine learning experts. It was the people that were just really good at crafting these prompts, just like spent a lot of time fiddling with it, playing with GPT-3 and so on. And so I think another thing that's happening here is just that it makes kind of really advanced AI accessible to way, way more people. And that's also what we've seen in terms of our user base. A ton of our users are building applications. They are not...As part of having played with GP-3, they start learning programming because they get ideas and they want to build companies around it rather than, you know, they come in as developers. And I think that's been quite powerful. And that

just kind of shows you that there's probably so much more things we can apply this to. And in some way, I think we're still really early here.

**Fernando Lucini [00:14:45]** Because we talk...we should talk a bit about democratization. I'm not quite sure you guys were aiming to democratize this kind of AI, but ultimately it solves the problem, which explains why software engineers or people like that are really interested because they want to solve problems, and this solves problems. And then it's about the gaining the skill like this gentleman - we should put in the notes for the podcast so people can go to his website because it's quite a fun one. So people that just become, as you say, expert in making them the machine do what they need. They understand what it likes, what it doesn't like. And we're talk anthropomorphizing a little bit more. I'm going to call it a by-product. You're going to correct me and tell me if it or isn't. But things like embeddings, so suddenly you're in a place where you have this model that can understand language and then we're trying to figure out how the world wants to use it. We're going to talk about use cases a little bit more, but we're effectively opening up to the world saying, world what would you do if the machine could help you with context and you know, and things like that? And some of the outcomes embeddings which for me as a data scientist and an engineer is the most fabulous things because it solves such a simple problem. I give you this, you tell me what's important and tell me what's important with a precision that I can actually count on to use in something else that I want to use over there. Right. How does that come out? I say I call it a by-product and I don't want to do that and get it wrong, because I have the utmost respect, I think it's the most useful thing. But how does that suddenly pop out of the machine?

**Peter Welinder [00:16:13]** Yeah, that's a great, great question. And actually, it was funny kind of because, you know, the thing that pretty much everybody with, you know, AI background, machine learning background asked us after we had released GPT-3 is like, hey guys, I want access to the embeddings because if I get the



embedding, I can build much more powerful things around this. And you know, it actually took us about a year before we released embeddings for GPT-3. And one of the big reasons was that it turned out that the straight embeddings that came out that you could get from these models... they were good, but they weren't great. So we actually had to put in a bunch of extra work to make these embeddings really, really, really powerful. But I think what we have found is that if you basically build embeddings on top of these large commercial models, like GPT-3 with a few extra kinds of tricks, then they have a much richer semantic representation of language. So, you know, what we've found is that these sorts of embeddings are much more robust to all kinds of semantic kind of noise and so on. What we've seen in practice is a ton of people now building on it. And I think you're right. You know, the way I look at it is sort of what is the modern kind of NLP stack becoming? And I think it's really kind of three things. It's like embeddings is a big piece because embeddings let you basically pull in relevant information and do all of these things like we like doing in machine learning. Like things like clustering and classification, you know, hooking it up to other systems like recommendation systems and stuff like that. But then I think the other big piece that GPT-3 enabled is kind of the, the more text generation piece. That was just the thing that just barely didn't work before as you know. Like even though it worked in some specialized cases like translation, but in most cases, it was kind of pretty, pretty bad, you know? And then the last piece is sort of fine tuning where you can take both of these kind of building blocks and make them even better with more data. But I think, for example, one area where embeddings become extremely useful is if you want to do a truthful question and answering, if you ask deeply three questions like, you know, what's the capital of the UK? It will know that. It would probably know that the Prime Minister is Boris Johnson. It would know all of these facts about the United Kingdom, but you can also ask it about kind of arbitrary topics, and chances are it will make something up. And that's definitely a pretty big shortcoming of GPT-3. Like the bigger these

models get, the more facts they know. But they don't have a great sense of when they know things and when they don't. And there's some really interesting extensions that people are building. For example, like DeepMind has a retro model that connects to its training data to become more truthful and so on. And the way you can kind of use embeddings, for example, is to find relevant sources, like give it to the model to be able to answer questions. So, for example, if you have a question about the United Kingdom, you can look through all of Wikipedia for things related to United Kingdom, but if you have embedded it all and find similar articles, give it to GPT-3 and GPT-3 can look at that information as it's answering questions and it can now answer it truthfully. So that's a way of like combining these building blocks now to you can get much better performance. And I think this is one thing that we'll see more and more. For example, in your example with emails like if you want to reply to emails, probably the best way automatically, for example, if you want to do that, to have like basically a bot that, you know, answers things as you would, you would call it similar emails, right? And you would give that to GPT-3, and it could probably do really well.

**Fernando Lucini [00:20:33]** Based on the fact that I don't read them or answer them which is what it's gotten down to. And we haven't spoken about it, but the world of search, which I also think it's one of those where embeddings and other things suddenly make more sense. By the way, we should throw a line so that it's really hard problem to solve, but because at the end of the day, you're trying to consume all that knowledge of a company and then you're trying to help people get what they want. But you're doing it with almost both arms behind your back because you're not context aware at all. To some degree, you're looking for keywords, and in some cases, you can get a little bit more advanced and do more interesting things. So suddenly you're in a place where, you know, context or comprehension is at the core of that, and you start doing the embeddings and you can start layering these things and with a bit of luck you get what you want. And again, it's a super





complicated problem where embeddings and other things are going to make a difference. And it just goes on and on and on the kind of things you can do, right?

**Peter Welinder [00:21:25]** Yeah. And let me say one thing about that, which is that, you know, the problem with search was always like search for the Internet worked well. Like the reason Google worked well was because of PageRank because it found all these kind of links between pages that gave them sort of a trust score and so on. And the problem everybody always had was once you went inside of an organization to do search, you didn't have all those links and you weren't able to build the same kind of rank and so on, and so all you had was keyword search. And we all know keyword search is not great at all. It's really brittle for synonyms and it's not at all able to like to match semantic concepts. And I think what these new, much richer embeddings are able to do is to have a deeper semantic understanding of text such that you don't actually need that sort of the PageRank stuff. You can actually find stuff that is much more relevant to a particular question you're asking rather than thinking about keywords. You can just like, you know, write the question, embed it, look at similar embeddings of all your documents and find the documents that are most relevant to answering that question. That's a thing that just was not possible, like even like five or ten years ago or it was really, really bad. And it's getting to a point where it's just like scary good today. And I think that.

**Fernando Lucini [00:22:49]** With the low barrier, which you're going to talk about, with a quite low barrier to entry, because the point of some of these problems are that when you don't have that comprehension, what you start doing is cooking queries, connecting things to things. It becomes very manual, and you can get to some degree of satisfaction. But you're effectively overcoming the fact that there's nothing here that has any comprehension of anything, right?

**Peter Welinder [00:23:09]** Yep.

**Fernando Lucini [00:23:10]** And we should ... when we talked about comprehension, and we should simplify that, maybe for the audience in simple ways. Right. And I know we talk about GPT-3 a lot because I'm fascinated by it, but you got things like you can ask questions. We've got something where we can ask the questions and you're going to get answers, which are going to be, you know, based on this enormous general knowledge basis. So you do that. You can interact in a natural language, I suspect, as natural as you're going to get.

**Peter Welinder [00:23:36]** Yep.

**Fernando Lucini [00:23:37]** So you can do that. These lovely examples of almost asking questions of GTP-3 around chess and just through the existence of games of chess in the data set to be able to literally forecast game place. I mean you can't make it up, right? You can ask if it finish the sequence and just through occurrences in the data and its understanding of those things for it to, you know, continue the sequence for you, right?

**Peter Welinder [00:24:05]** Yep. You know.

**Fernando Lucini [00:24:07]** And I guess I'm saying all this because you could go on and on and on and the limit is what needs do you have from the comprehension of the data, right?

**Peter Welinder [00:24:19]** That's right. And you know, the crazy thing about this is I think, you know, GTP-3 has been around for about two years now. And I think there's still so much work that needs to be done in order to understand the limits of this technology. With the models we already have today, there's so much more things you can do with them that I don't think we've fully discovered yet. And let me give you one example of this, and this is an example of how you get the models to kind answer questions better in some way and if you have hard questions the way a human answer them, they just they don't answer just impulsively immediately. They think about it a little bit. You know, they do a bit of reasoning in their head.



And it turns out that you can actually get GPT-3 to do the same thing. So there's really interesting research, a very famous kind of data set in this domain, it's a data set called grade school math where you do things like you ask kind of simple math questions, but it's questions that need a little bit of reasoning. For example, if I ask you a question like, you plant a tree next year and then you plant two times that tree in the next year, and then you plant two times more trees the third year or so.... you're kind of doubling the number of trees. But then every other year some deer comes in and eats some of these saplings like 50% of the saplings or like 30% of the saplings, you know? How many trees will you have ten years from now? Like, you know, I don't think you will be able to answer it.

**Fernando Lucini [00:26:16]** And I was going to say.... So I've got a 14-year-old autistic child who would be unable to answer this question because his mind is so extremely logical that would be like 'that doesn't happen in the real world. What are you on about? So it's overlaying the world of logic with the world, as you say, of reasoning, and reasoning is complex, nuanced. It requires understanding. It requires your knowledge of the world. Right. So that's why it's tough, right? That's why my little Jacob would sit there, and a look and you go, what are you on about? It doesn't happen in the real world because his world doesn't work like that. It's a beautiful mind, but it just doesn't work that way, right?

**Peter Welinder [00:26:56]** Yeah, exactly. But what's really interesting is that if you ask a human this sort of question, like even if they cannot answer it immediately, they can sit down with pen and paper, like work out kind of a solution to it. Or like, you know, if they're really good, they can reason through in their head. But, you know, it's going to take you them a minute to kind of do it. Right. And it turns out that, you know, you can ask GTP-3 to answer questions like that. And, you know, it would be wrong most of the time. But if you tell it to first reason through it, like, you know, go through the steps,

like how would you solve this problem? And they kind of go and kind of keep track of things, then it can actually solve these sorts of problems. And I think this is really interesting because you can do the same thing for like questions. Like if you have a really complicated question on history, for example, you might want to do like a comparative analysis, and this answer to the question is a summary of that analysis and so on. And I think this is what we will actually see kind of more and more of is just like where people are really good at teaching GPT-3 to do the reasoning that it needs to do in order to solve quite complex problems. And it turns out that the models we have today are already quite good in this domain. It's just that we haven't really pushed them far enough yet.

**Fernando Lucini [00:28:14]** Thank you very much for listening to Part one of our recording and we look forward to seeing Peter again in our Part two.



and implement something that could be very impactful for their organization workplace. And with that I would like to share my sincere thanks to both of you for accepting an offer to speak to us and wishing you all the best.

Copyright © 2021 Accenture  
All rights reserved.

Accenture and its logo  
are registered trademarks  
of Accenture.x