

# The Just Word Podcast

## Transcript of Episode 36: The Future of A.I.

Guest: Peter Voss

**Pat Bolland** 00:00

Peter, a real pleasure to meet you.

**Peter Voss** 00:02

Yeah, thanks for having me.

**Pat Bolland** 00:04

OK Peter, let's start with the real basics, because I don't understand sometimes what's going on with artificial intelligence. And you also have introduced another term called artificial general intelligence. So I need to know the difference between those two, if you don't mind.

**Peter Voss** 00:20

Yes, absolutely. The term artificial intelligence is actually quite confusing. Because if you go back to the original meaning of AI, you know, when the term was coined some 60 odd years ago, it was really about building machines that can think and learn and reason the way humans do, you know, the kind of AI you see in the movies, you know, where you can actually have a full conversation, understands, has its own motivation ... you can, it can do jobs for you. And unfortunately, in the movies, often it then tries to do bad things, you know. But, so that was the original meaning of AI. And 60 years ago, they thought they could build machines like that, you know, within a few years. Now, it turned out to be much, much, much harder than that. So what happened over the years, the field of AI has really changed into solving very specific narrow problems. So instead of having this general, human like AI, your or intelligence, you have machines, that sort of programs that solve very specific problems, narrow, very narrowly. So a good example, here is when IBM built deep blue, which became the world chess champion, you know, it was a specific system built to be very good at playing chess. Now, that's all it can do. It can't play checkers, it can't, you know, do trip planning or do anything at all

**Pat Bolland** 01:52

Right. So is that what we are ... sorry to interrupt Peter, but is that what we would now call machine learning?

**Peter Voss** 01:58

Well it's morphed into ... machine learning is really the part that's exploded over the last 10 years or so. And I'll talk about that in a moment. But yes, exactly. So it's solving one particular problem, but there is, there's actually a deeper issue involved here. And that is that it's the intelligence actually resides in the programmer, much more than in the program. So the original idea of AI was that the program has the intelligence to, you know, learn how to play chess and solve other problems and you know, like a human basically learn new skills and solve problems. Whereas with narrow AI, the program is figured

out, how can we use computer technology to solve one particular problem? So that's narrow bias creeps in

**Pat Bolland** 02:51

So what you're saying is garbage in, garbage out kind of,

**Peter Voss** 02:54

Well, that too, but you have that problem, you really have that problem for people as well. So, now, you talk about machine learning. Over the last 10 years, company, companies figured out how they could use massive amounts of data that that accumulated a massive amounts of computing power, you know, like, obviously, Google and Amazon, how to use that, to build models that can solve computer models that could solve particular problems, such as, as image recognition, you know, how to recognize a person or a motor car, or something, which is, of course, used in autonomous cars, or how to recognize speech. And it has become incredibly successful in a number of areas. But again, it's narrow AI, it's solving one particular problem, one particular class of problems. And a lot of that intelligence, again, is external is in the data scientist in in figuring out what data to collect, how to mock up the data, how to specifically train it. And you know, it's then using the computer technology to solve that particular problem. So, again, if you have, you know, a system that can become the world champion and go, if you change the rules of the game, it can't cope with that, you know, it wouldn't. Again, it wouldn't be able to do anything that it hasn't been specifically trained on. So it's this trend of going towards narrow AI that's been very successful in solving a number of problems like targeted advertising to, you know, the bane to all of us. But of course, incredibly a lucrative, you know, for companies developing sort of targeted advertising using using machine learning. But enough number of us a small number of people actually want to recapture the original ambition and the original dream of AI to build Thinking Machines that actually have the intelligence inside them, to learn new things, to solve new problems, and to be able to be able to think and reason like humans. So in 2001, I actually got together with a number of other people with that idea. And we coined the term artificial general intelligence or AGI, to basically talk about that type of general AI. And that term has actually become quite generally accepted, surprisingly, but it's still a very small number of people actually working on what you could also look at the real AI, you know, or the, the original AI.

**Pat Bolland** 05:57

Okay? How is AI different than your, or my, brain?

**Peter Voss** 06:06

Well, from a brain point of view, it's very, very different. Of course, our brains are, you know, a biological material built from biological material. And, you know, designed if we can use a word designed by evolution, you know, formed by evolution of millions of years. Whereas an artificial brain uses computer technology, and its engineers who are building it. So it's really very, very different. Now, having said that, there is definitely inspiration and ideas we get from the way our brains solve intelligence. You know, and so even to the terminology that people now talk about neural networks, you know, like neurons could be in connected together. And so that's a useful model. And in fact, what we do is we have these artificial neural networks. But realistically, they are very, very different from the way neurons work in our brains. So it's really more understanding what intelligence is, and what thinking is

really more how our mind works. So you want an artificial mind, you want a mind that can think and reason and learn the way humans do, but done in a very, very different way. Now, I think, a reasonable analogy here is flight. You know, we've had flying machines for over 100 years now. We are still nowhere near reverse engineering a bird, you know, so you could say, you know, are we building flying machines that are like a, like birds? Well, not really. But it's a flight part we want to achieve. And we've been successful in doing that. But not in the same way that birds do it not the same way that, again, evolution and biology of solve the problem of flight. So in that same way, we want to solve the problem of intelligence or the requirements of intelligence, not in the way, nature did it. But you know, using the tools and materials that we have available to us.

**Pat Bolland** 08:18

You know, it's funny, you mentioned that the early iterations of flight were Leonardo da Vinci, and he actually did try to redesign the bird. But obviously, it didn't work out. And we use a more structured vehicle as it stands currently, let's get to a practical component, you know, I've driven a Tesla, for instance, and it would represent I think machine learning, it was able in a rush hour situation to change lanes, for instance, and this was a few years ago. So this is old technology. But what we run into on a more ubiquitous basis, I think a lot of people is when they go to sign into access their account at some kind of a service, they get a thing on the side called a chatbot. And these things seem to be more and more intelligent. Would you consider those artificial intelligence?

**Peter Voss** 09:07

Yes, certainly. And so again, there are distinctions to be made. First of all, when you have a lot of windows windows pop up, you may actually be talking to a human. So typically, if, if you have a really good experience, and there's a good chance there's a human actually giving you the answers on the other side, but, you know, they again, there's sort of the the narrow AI approach, which is, you know, keyword spotting, you know, what, what keywords are you using, and, you know, you can see the same thing in Alexa and Siri. You know, if, for example, I say, don't ever give me Uber again, you know, I hate Uber. Don't ever give me Uber again. It'll probably pop up the Uber app because it recognizes Uber and says, Well, that's probably exactly what you don't want, right? So you know, the is not a real deep understanding. So a lot of the chat bots and a lot of frustrations that people have with chat bots are that they're not intelligent, they don't really understand you deeply. It's just like keyword spotting. And then they put you through some kind of like step by step flowchart process, you know, ask you a question, yes or no, enter this and so on. But anything that wasn't anticipated, it goes off the rails very quickly, you know. And that's because they really don't understand. They don't have memory, they don't remember what was said earlier. So that's a typical chatbot, you have, and we, you know, what we are introducing is a chatbot with a brain. So that there is actually deep understanding that it remembers what you said earlier in the conversation, or even in a previous conversation, that it can apply some common sense to what the conversation is about or what you want, and can reason about things. So that's kind of the, the upgrade you're really hoping for, is to have a chatbot with a brain as opposed to just, you know, a chatbot, that just as key word recognition, and takes you through a flowchart. Now, having said that, there are applications if you just want your bank balance, you know, sure. You don't, you don't need a lot of sophistication, but then people will typically just go to the website anyway, to get that, you know, it's usually more complex things you want to do when you when you activate a chatbot.

**Pat Bolland** 11:31

Okay. And that's what you do at A.I.go?

**Peter Voss** 11:34

Yes. Yeah, that's, that's what we do.

**Pat Bolland** 11:36

Okay. Against that background? How will the client experience change? Do you think, if you get your way?

**Peter Voss** 11:45

Well, you know, we talked to a lot of large companies, and most of them are actually very disappointed and frustrated with their chatbot implementations that they have, because, you know, they don't have good customer uptake. You know, one of the biggest banks, you know, have like a 2% uptake of a system that I developed for 10s of millions of dollars, you know, or even hundreds, because the system really doesn't have a brain, you know, it doesn't understand. So, I mean, imagine you're talking to a human, or you have a personal assistant, and they didn't remember what you said two sentences ago, you know, or they just didn't pay attention and just kind of took key words that they hear, you know, and try and answer something, that that wouldn't be very satisfying. So, what customers really want, and therefore the company's providing service is, you know, something that can be hyper personalized, they can have deep understanding, and also remember you, and, you know, one of our clients is one 800 flowers. And what they, what they were looking for, is to put a chatbot in, that's more like a concierge that knows that each individual, it knows what your preferences are, who you buy gifts for what kinds of things you like, and, you know, people you buy them for, what they like, what the status is of your, your various orders, and so on. And it can proactively, it can talk to you like a service agent who really knows about you, in fact, better than the service agent, because it will have perfect memory and perfect recall, of, you know, what your situation is. So that is I think what customers ultimately would would like is to have this hyper personalized assistant, you know, whether that's for help with your cable service, you know, then you're troubleshooting it. And it might recommend, well, you know, reboot your, your, your router, you know, now, next time you call in or you're contacted, it wouldn't ask you the same thing again, you know, like, we have now, every time you call in well, have you tried rebooting the router? You know? Yes, I've tried that three times. Okay, why don't you try and move it into the kitchen, you know, see if that'll work. And then you'll call back and say, Well, did that work? So having that hyper personalized service, I think is really ultimately what, what people want and the company so that serving them,

**Pat Bolland** 14:21

You know, Peter, in my mind's eye, what I'm seeing is the movies and Ironman and Tony Stark talking into his helmet and I forget the name of the butler that he has there. One area that I thought would be perfect for AI and I've been disappointed it's wrong word but surprised that it hasn't moved into is finances and investing. How do you see that playing out over time?

**Peter Voss** 14:49

Yes. It's the same. It's the same thing as again having a chatbot with a brain, you know, that can really understand and reason and have common intent and no know what it's doing. So it's, that's really what we're looking for going forward. Now, financial institutions obviously have the additional requirements of, you know, security and liability that they might have if the system gives you bad advice or wrong advice, or whatever. So there are those additional considerations. But it's still having that hyper personalized service that that it's a perfect application for that, that knows what kind of investments you currently have, what your profile is the kind of, you know, risk profile, you have all of those kinds of things, your your prior transactions, taking that into account. And incidentally, what, what we often find, and research has found is that there a lot of areas that people actually prefer to talk to your computer, if it can do the job, then to talk to a human, because there might be, you know, it might be something they embarrassed about, or they don't want to waste a human's time, you know, talking endlessly with something they may not understand very well. And they don't want to bother somebody with it, you know, so or as I say, they might just not feel comfortable talking to people, obviously, medical conditions are like, Yeah, but I think financial situations that you're also fall into that category. So absolutely, we will, we will see, I think, significant improvements going forward, as we have more chatbots with the brain.

**Pat Bolland** 16:35

It seems to me that there are limitless applications for artificial intelligent intelligence AGI in your particular case, what's missing, as it stands currently, and how quickly is the space evolving?

**Peter Voss** 16:52

So from my perspective, what is missing is that not enough people are working on AGI. And it's sort of an accident of history in a way that deep learning machine learning has become so incredibly successful over a short period of time, because there are a few large companies that have massive amounts of data, and, you know, massive amounts of computing power, they are now building models of, you know, literally hundreds of billions or trillion words and parameters that they get from, you know, everything from internet, Wikipedia, and so on. And they built these massive statistical models that can actually have very impressive conversations. That sounds that sound very intelligent. But when you scratch, below the surface, there's there's nothing there. There's no, there's no understanding, there can't actually be used for things like we were talking about for customer service or so. But because it's been so impressive in areas such as targeted advertising, and image recognition, and so on, all the money is flowing into that area into this narrow AI, deep learning machine learning. So I think the biggest advances we'll see, as more focus shifts to AGI were people saying, We really need the intelligence in the system, and not the person building the model, you know, solving the particular problems that we don't solve one problem at a time with external intelligence, but that we solve the problem of building a machine that has true intelligence internally. As soon as we get more money flowing into that at the moment, you know, less than 1% of all the AI funding goes into that.

**Pat Bolland** 18:49

Okay, final question, then, Peter, if AI does advance, how do we stop it from being bad? As you said at the beginning, the movies always have it as bad?

**Peter Voss** 18:58

Well, it's unfortunate, but you know, obviously, it makes, supposedly, for a better movie. I was asked once to consult on an AI movie. And, you know, I went through the script and made certain suggestions. And I said, but wouldn't this be a much better story? If, if the ending was like humans and AI's work together? And, you know, we had utopia with it, and now they weren't interested in that. That movie wouldn't sell. So I think there's a big gap between the perception of AI has to be bad when it becomes really intelligent with the reality. You know, AIS aren't going to have that reptile brain that humans suffer from, you know, that we have problems with, with with ego and, you know, of course, evolution designed us to, to focus on reproduction and survival. You know, that's what is very inherent, of course, in our sort of, animal nature. I'm and our rationality, civilize us, essentially. Now an AI isn't going to isn't going to have that starting point of being built to survive and reproduce, it's going to be fundamentally built to be intelligent and to do things for human. That will be its, you know, DNA, if you want to call it that, that will be its purpose and its DNA. So I don't have any reason to believe that an AI will, you know, have any other motivation than that. Now, it's a very complex topic. Of course, you know, you we have actually written quite extensively about issues of freewill and about consciousness and ethics, you know, how do we know right from wrong? And how will computers know right from wrong? And I think there's actually some pretty good answers and in my mind, the old point towards AIS AGI is helping us to improve our lives significantly.

**Pat Bolland** 21:05

We can have a discussion forever on the reptilian mind. Great descriptor. Peter, thank you so much for your time.

**Peter Voss** 21:13

Yeah, thank you. This was fine.

**Pat Bolland** 21:15

Peter Ross ...He's the CEO and Chief Scientist at Algo.ai