

AI and You

Transcript

Guest: David Wood

Episode 20

First Aired: Monday, November 2, 2020

Welcome to episode 20. Today's guest is David Wood, one of the pioneers of the smartphone industry, co-founding Symbian in 1998. He is now an independent futurist consultant, speaker and writer. As the Chair of the London Futurists, he has hosted over 200 public discussions about technoprogressive topics. He is the author or lead editor of nine books, including *Smartphones for All*, *The Abolition of Aging*, *Transcending Politics*, and *Sustainable Superabundance*.

I start out the first half of our interview talking about a possible shared sense of a rendezvous with destiny, and one of David's responses to that was to say that there is no predestination, which I very much agree with, because we need to know that we are architects of our own future and we should take command of our direction. For those of you sci-fi buffs who are fans of Isaac Asimov's *Foundation* series – like I am – there is no Hari Seldon or psychohistory predicting our every move for the next thousand years.

What I was getting at when I was talking about the shared sense of destiny was a growing feeling that the human race is on a course that leads to something momentous, and it may either be a future that is very welcome or one that is disastrous. That's embodied in the tagline of my book *Crisis of Control*: "We're heading for utopia or Armageddon... what can you do?" It's also in line with a cyclic theory of history I find very credible and have mentioned before, from the historians William Strauss and Neil Howe, that predicts we are in an era characterized by momentous crises and apocalyptic responses. You can read more about it in their book called *Generations*. If you also feel that we're accelerating towards something big, you might feel encouraged by their prediction that once this era is complete, we'll enter a phase of good feelings and prosperity. As long as we make it there in one piece.

David mentions – or he will - the concept of the *Singularity*, and that bears some explanation. The idea of the Singularity is that exponential technological progress will accelerate to become so fast that we perceive it as infinite. The word was first used in this context by the mathematician Stanislaw Ulam in 1958, when he wrote of a conversation with John von Neumann concerning the "ever accelerating progress of technology and changes in the mode of human life, which gives the appearance of approaching some essential singularity in the history of the race beyond which human affairs, as we know them, could not continue." It was British cryptologist Irving J. Good (who happened to have been Alan Turing's statistical clerk during WWII) who described what this meant in 1965 when he said "Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an 'intelligence explosion,' and the intelligence of man would be left far behind. Thus the first ultraintelligent machine is the last invention that man need ever make."

But the real popularization of the term Singularity has come from Ray Kurzweil, in his book *The Singularity is Near*. (By the way, shortly he will publish the sequel, *The Singularity is Nearer*.) Kurzweil also started the Singularity University, which I'm a member of. Now the word singularity has a particular meaning in mathematics, but it is best known in cosmology, where it's how we describe the point at the center of a black hole where everything goes to infinity and the laws of physics break down. So that makes for an arresting image when we apply it to human progress, but it has problems. A black hole has a very sharply defined edge, so this suggests that the human singularity would too. But does that mean that the Singularity is going to happen at, say, 6:17 pm on October 23 2045? That prompts questions that probably aren't very useful, like, what will things be like 5 minutes before that happens? Once we get past obstacles like that, if you accept that a technological acceleration that outstrips our adaptive abilities is possible, then if you think it is likely to happen in less than massive time frames and that we should make sure that humans survive and thrive through it, then that makes you a singularitarian, which is how David describes himself.

Which is a good point to go on with part one of my interview with David Wood.

Hello, everyone, I am here with David Wood of the London Futurists, one futurist to another. David, welcome to the show.

Peter, it's a real pleasure to be here.

And I've been reading some of your output here, and we are very much in alignment here. I am really fascinated by the phenomenon of how many people are around the world, coming to the same conclusions at the same time. It's almost as if we are on a rendezvous with destiny here in the human race. What's your sense of the future, as a futurist? What's your sense of the velocity with which it's arriving?

Things are changing more quickly than previously people thought possible, for all kinds of reasons. Of course, it's the connectivity that we have but it's also the way in which we can see other people changing. And in the past, people were reserved because they thought, "Well, other people like me won't change quickly." But now we can see lots of other people visibly changing their attitude, which encourages us to change our attitude too. It used to be, for example, that if you asked whites in a neighborhood, would they mind if people from different racial groups moved in, they would say, "No, I don't mind." But if you ask that same person, "What do you think your neighbors would think?", they would say, "Well, probably my neighbors wouldn't like it, and they could be disruptive." Now, that kind of misunderstanding is less often and it's part of why ideas change more rapidly, but not just for good. Ideas sadly are changing rapidly for bad too. We are vulnerable, not just to biological viruses, we are vulnerable to psychological viruses, to various kinds of excitement, which are actually bad for us psychologically and socially, rather than good for us. So the pace of change is increasing. You mentioned at the beginning, you thought there's a kind of date with destiny - more people reaching the same view. Yes, but also there are lots of people taking very different conclusions. And as a futurist, I want to tell people clearly there is no predetermined destiny. There is no inevitability. I reject the idea of technological determinism. I think it very much does matter what we humans do. I say, individuals and as groups to steer technological development to ensure that we think through the

implications beforehand and we don't just jump onto some train that looks like it's going to an exciting place, but by the time we arrive, we think, "We should have paid more attention."

I couldn't agree more. In fact, I like to say that the reason that there are these diverse, diametrically opposed views of the future, so much now that people talk about artificial intelligence as creating a utopia and other people talk about it as the end of the human race, and why are these as far apart as you can get? I think that it's because we haven't decided yet how we collectively are going to grow up to use this tool. And being handed that kind of power is like giving a toddler a chainsaw. Yeah, they might cut down a tree, but they're more likely to cut off their leg. So it's time to grow up. Let's back off a minute and just get into your history as a futurist. I'd like to get inside your head and find out how you came to get into the field, how you came to give yourself that label. What's something of your personal history that led you to decide this is what is the most important thing for you to do?

So I spent 25 years in the mobile computing and smartphone industry from around 1988 when I got a job as a software engineer in a company making laptop computers, and eventually handheld computers. For a while, I was a software engineer, but as I moved up the management hierarchy, I was very much involved with colleagues in trying to say what the future would hold. What will people want to do with these devices in a few years' time? What will the technological possibilities be that our software and hardware need to anticipate? And eventually, from 2004, I became the executive responsible for research. It was my full-time job to understand the trends, and to try and alert the rest of the company and the wider ecosystem to not just incremental developments - we will pretty good already at responding to incremental development - but it's spotting disruptions, which is something that initially seems uninteresting, initially, it seems of less importance but as it develops, and as it combines with other trends, often it can turn the world upside down. So that was my role. And when my company, Symbian, the creator of a successful smartphone operating system from the noughties when we went open source in 2009, in part to try and respond to the open-source Android from Google, we decided to change our job titles to be a little less corporate and a bit more peppy. And so, we threw away "Executive Vice President for Research" and I became "Futurist and Catalyst", or "Catalyst and Futurist". And I was talking not just about changes in how people develop and use products, but more broadly, the wider social implications of new technology, including artificial intelligence. And eventually, when I went independent, from around 2012, I have been more or less full-time independent. I've been marketing myself as a futurist and catalyst and author. Sometimes I slip in the word singularitarian as well, just to say, "Hey, I take the singularity concept very seriously." It is the most important thing, all the distractions, all the important things, let's say, of the present, will be put into shade if we're not careful by the increasing power. And as you said, it's a bit like a chainsaw in the hands of people who are not trained to use it. Wow, it can be powerful but oh, my goodness, it can be disastrous, and as somebody who has taken possession of a chainsaw in the last few months, I know that it's easy to be distracted and bad things can happen if you let your mind wander. And what we have now with AI, but also with growth of nanotechnology, also with biotech possibilities to reprogram ourselves genetically, epigenetically at the biome level, and so forth. And last but not least, also with technologies to enhance our brains,

cognotech. It is much more than a normal chainsaw and we all need to pay attention to it, which is why I work as a professional, but also why I work in various public groups such as London Futurists to try and raise the caliber of the conversation. So that instead of just having a kind of giggle about technological possibilities - the future is entertaining, of course, but we also take it seriously. And we try to find ways to help people to distinguish between science fiction and science fact, and between the scenarios, which initially are attractive, but which will turn out to be bad for us and vice versa. Some of the initial scenarios that make us go “Yuck”, or in retrospect, when we figure out different ways to approach them, could actually be very good for humanity. And we need to help people to improve their ability to have that kind of conversation to go beyond that Future Shock, which is unhelpful.

I very much appreciate the commitment that you're evidencing here to raising consciousness around this. It's frustrated me that I could almost certainly make more money telling people how to lose 10 pounds in three weeks than I could telling them how to prevent the end of the world or deal with an oncoming utopia, or how they could act to change the course of whether we get one or the other. That's one of those paradoxes. And people say, “Well, you can't change human nature,” but I have this sense that we may have to in the future. But perhaps that's another conversation. You are an author of several books: *The Abolition of Aging*, *Transcending Politics*, *Sustainable Superabundance*, and *Raft 2035*. You're really putting it out there in the communication to people about the future and as you say, not just entertaining, but engaging people like the London Futurists and trying to make a practical difference, and that's something I'm very much working to do myself. This podcast and my books are supposed to be “What can you do?”, not just “Let's be a spectator.” The future is not a spectator sport. How do you use groups like the London Futurists to advance this beyond the level of a spectator sport?

So part of it is about pointing out interim milestones, or sometimes called canary signals, using the idea that when humans used to work in mines, the humans would sometimes not notice that the atmosphere was changing, whereas as a small bird, the canary, would fall over on its pouch in a small cage it was brought down. So it's attention to smaller things happening, which could be signals that larger changes are happening, and are possible. So some of the events are focusing not just on how this technology might work out in 10 years' time, but how it might work out in 18 months as well, in terms of things that could change people's lives now. So we had a meeting, for example, on implants. Implants are a step beyond handheld computers, a step beyond wearables, they are chips that can go inside the body. And there are some geeky people who put an implant in the web, in the skin on their hands between their thumb and their forefinger, but that's only the start of what's possible. And so that's one example of a meeting when we looked at the roadmap of a company that had big long-term visions, but also had things in the short-term, which made people think, actually, “This isn't just a geeky thing, this could improve their health.” And this could be beneficial to keeping track of animals or even dare I say it, elderly grandparents, who might be wandering off without anybody knowing where they are. So the practical discussions take place from there. But no, it's hard to make business people really focus on this because although individual business people may be curious, may be

fascinated, the way businesses are set up is to focus on what can happen in the short-term, which is why we have to work harder to indicate the ways in which the longer term potential already can be prefigured in the shorter term developments. Which is why I talk a lot about “Here are ways in which artificial intelligence could be different from today's in just five or ten years’ time.” And so companies better start thinking about paying attention to these possible changes in AI. It's not yet that we will have artificial general intelligence, probably not. That's probably further down the road but there are ways in which artificial intelligence is changing. So it's not just a matter of applying best practice, which is already here, and that's difficult enough, of course, but it's also a matter of figuring out some incremental, small disruptive changes, which might change the landscape of AI in just a few years’ time.

And, to some extent also, we have to prepare people for the unexpected or the unexpectable - the things that you can't anticipate. In 2005, I don't know how many people could have foreseen how much Deep Learning would have changed the landscape of computing five years later. And so that was an evolution there in computing that changed artificial intelligence, the shape of it. And to think that we're done with that kind of evolution, because I see now people making linear extrapolations of what Deep Learning can do, and yet, there must be limitations to what it can do. But to think that there isn't another breakthrough like that coming in five to 10 years would surely be blinkered.

Exactly so. In fact, even by 2010, there was very little perception of how much Deep Learning would change things in just a few years’ time. There were, of course, some people working on Deep Learning, but they were very much the minority. They were viewed as strange by the mainstream of AI specialists. And if I look at my own slides from 2010, which I still have on my laptop - I've got slides going back to 1992 shockingly - and look at the very first PowerPoints I did then so I can see the evolution of my own ideas, and I think it's a useful discipline to look back and see what you foresaw and what you didn't foresee. So in 2010, I was very excited about Google Glass and the wearable computers coming, and there was nothing much about Deep Learning there. There was a lot about AIs going to get smarter, but not in the particular way of neural networks, which caught people by surprise because people thought there were limits, but it turned out that through clever tweaks to the algorithms, backpropagation, and via larger sets of labeled data, such as the image net that Fei-Fei Li and others at Stanford created, lots of images labeled. And also, unexpectedly, as it turns out via different kinds of processor, the GPU, which was developed for a different purpose. It was developed for graphics capabilities in games. And it turns out that that parallel processing in GPUs was also suited to doing lots of matrix multiplications at the same time. So in just a few years, from 2010 to 2012, 13, 14, 15, suddenly, the strange people who were talking about neural networks were the heroes of artificial intelligence. And we looked back already at what we started to call good old-fashioned AI. Now, I think, as you've mentioned, it's quite possible that in another five to 10 years’ time, we might look back at today's stuff and say, “Oh, that was quaint, wasn't it? Deep neural networks.” We got quite a lot done, that we didn't realize that we needed, and then what is it? Might be one-shot learning and the abilities to do not just learning from big data, but from small data with transfer learning. It might be things to do with all the evolutionary algorithms, which have been around for a long time, but like the neural networks never seem to gain much traction. And it might be

new insights from studying the human brain, whether it's software insights, such as Jeff Hinton's ideas on capsules, which have gone through various generations, or it might be copying other aspects of the hardware in neuromorphic computing. Or it might be things to do with generative adversarial networks, which is a kind, of course of deep neural networks but as that goes, suddenly, things happen, which nobody expected. And then who knows what's going to come out of GPT version three, four, and five that OpenAI has done, and now many other companies are saying, "Hey, we can do that too and we can do it better." And as that evolves, it might leave behind some of its embedding in deep neural networks and other features may come to be dominant. So now we can't know. And this is the uncertainty you refer to. We can't know which of these tracks are going to be the ones that are most significant hence, the greater importance not just of foresight, but the greater importance of agility and resilience. We should no longer optimize our systems for ultra-efficiency for just in time, we must re-engineer our systems, both at the corporate level and at the wider political level, to have more slack and to be able to change things more rapidly, rather than being bogged down in lengthy processes because we will be surprised. Some things in retrospect will be foreseeable, but other things, we had no reasonable way of knowing that would happen. So as a futurist, I'm trying to help people study trends more thoughtfully, help people to imagine the ways trends may intersect and transform each other, creating something brand new, but also, I'm advocating agility and resilience. And a lot of my own professional consulting, it's not just about transporting, people invite me back to say, "You talked about this agility thing, how does it work in practice? How do we make sense out of uncertainty when so many things are unsure?" And then we talk in such engagements about short-term sprints, and failing forwards, and about smart pivots, where you can benefit from what worked well, and also learned from what didn't work in that. And companies who have got the ability to do this agile learning will be the ones that are most likely to adopt quickly whatever new strains of AI, whatever new strains of decentralized systems emerge, and other new practices. Being expert in machines, doing the same thing again and again, minimizing variation, there's so-called Six Sigma approach, that made sense when things were less changing, when things were the same again and again. But that discipline is no longer anything like as central as it was in the past.

I, again, agree. I think the challenge is not to build a better machine because that's going to happen anyway. There are enough people working on that. We will get that. We don't have to push on that. The challenge is to build a better person to deal with that, and because the machine evolution could outstrip our own. You mentioned agile. I remember back in 2000 and I was meeting with my editor at the time because we were both at a conference in Monterrey. And he said, "Have you heard of this thing called agile development?" And I said, "No." And he slid a book across the desk to me. He said, "We've just published this." And it was Kent Beck's book on agile development.

Oh right.

He said, "We can't print this fast enough. It's flying off the shelves."

Extreme Programming was that, right? [By] Kent Beck?

Yes, *Extreme Programming*. Yes. And that was clearly a philosophy that could apply to many things other than just software development. So you now start to see this expansion of it into other fields of agile this that or the other. Wow, so many threads to pull on. Let me pull on one that you mentioned a little while ago, SingularityNET. I'll explain in the preamble to the show what Kurzweil's singularity is. Tell me about SingularityNET.

Well, SingularityNET is one of the companies that might play a key role in advancing AI in a quantum leap, rather than just an incremental way. So the insight behind SingularityNET, which is the brainchild of Ben Goertzel, is that you can reach a higher level of intelligence by allowing modules of smaller capability to network together, and they will talk to each other in various APIs. And what SingularityNET is providing is sort of an API for APIs to allow different modules to find each other, to exchange requests with each other, sometimes to offer various tokens and payment for requests being fulfilled, and then to provide services in new ways. It's probably what happens in the human brain. We're not sure yet, of course, but there probably isn't a single master neuron or a group of neurons. It's a network architecture of the brain that provides us with our sense of consciousness and general intelligence. So SingularityNET is a way of designing the APIs in which lots of these different modules might interact. And the vision of the SingularityNET is to enable benign AI. And one reason it will be different from what some of the corporations are working on, is with the focus on open source as far as possible, keeping things as transparent as far as possible. Whereas most AI today is being developed to help what? Spy on people, influence people's behaviors, sometimes to kill people or hunt them down in drones various times to make us more excited or angry in how we interact with social network, how we should gamble more. And we need the technology that is being developed less for these purposes because we don't want AI to grow up having that kind of education in a sense. We want it instead to grow up with a sense of open collaboration in which the motives, the goals, the structures of the AI are accessible for humans to look at, and indeed accessible for other AIs to look at so there can be more of an awareness of the possible failure modes of our AIs rather than them hiding behind some black boxes. So that's part of the vision for SingularityNET, and from time to time, I do offer my advice to Ben Goertzel and others in the leadership team there.

So it's like creating an electronic ecosystem for AIs to be embedded in so that - I'm paraphrasing here, but tell me if I got this right - so that instead of point solutions here and there that are - no matter how evolved they are like Google Duplex say, but that's just operating in its own silos, as is Amazon's Alexa - but the idea is to create a framework like a protocol stack for AI to interoperate with other AIs, is that the idea?

That's right, yes. And it's similar to what happened in my own background in - if I refer again to - the smartphone space. What's made the smartphone so interesting, isn't just the functionality that is built into them. It's the wide ecosystem of developers who come up with all kinds of unpredictable games and apps of all sorts, which were not in the minds of the designers of the smartphones, but which it turns out has great interest. So one of the first popular apps for the iPhone and the iPod was something that they would keep the very young children engaged. And if you'd said in advance, "Hey, this is what these devices will be used for", you might have been criticized as being a terrible parent for imagining that you could have some of your parenting

responsibilities delegated to a screen. But that's what happened, and many parents are glad, at least sometimes for that kind of app. And there are many other things which were not foretold. Again, if I go back to the predictions which I and my colleagues made, we did talk about implications for health care and education and navigation but many of the ways in which social media have worked, both for bad and for good, weren't really in our minds at that time, I can say. Well, let's have an open ecosystem for AI modules as well and that will allow competition to take place with many more people than can work just in the large AI companies. As I think it was Bill Joy from Sun, one of the Co-founders said, "Most of the smart people in the world do not work for your company." So how do we enable the smart people outside to contribute to our products? It's with openness and transparency - the right kind of APIs.

That's the end of part one of the interview.

I think it's fascinating to think about the relationship between superintelligent AI and the singularity. My point of view is that they will both happen at the same time. I do not think the singularity happens without artificial superintelligence, because you need a recursively self-improving computer to make exponential progress that isn't constrained by human limitations. Conversely, if you have an infinite rate of technological progress then it will by definition create everything that is possible, and that will include artificial superintelligence.

In our ripped-from-the-headlines segment, General Motors submitted a petition to the National Highway Traffic Safety Administration to grant them permission to deploy GM's *Origin* self-driving car. Now, most self-driving car laws are made at the state level, but the reason GM needs the Fed's OK is that the *Origin* has no steering wheel or pedals. This isn't a concept car; the *Origin* is planned for production at GM's newly renamed Factory Zero in Detroit.

Now this is arguably overreaching. We're not close to having autonomous vehicles reaching level four or level five capability where they're able to do in particular the last mile in urban or rural environments where there are just so many orders of magnitude difference between what they have to deal with there versus what's necessary to handle on a freeway. I think we've doing ourselves a disservice, in fact, in creating the five levels of automotive vehicle autonomy that is part of the Society of Automotive Engineers standards, because we think that it must be as much of a jump from level 3 to 4 as it was from level 2 to 3.

In fact I think it would be more like a thousand times the distance, if that. And it would be more helpful to think that we're currently at level 3.010 and advance to level 3.011. That aside, General Motors is going ahead with this *Origin* car anyway; let's see what happens.

Now, Norway is now deploying ferries that are self-driving – or self-sailing. The small electric ferries "work like an elevator" according to Zeabuz CEO Erik Dyrkoren. They take passengers across a canal in a 60-second journey that saves them a 15-minute walk.

Next week we'll conclude the interview with David and talk more about the implications of SingularityNet, how we should adapt our economic systems to the abundance of artificial intelligence, the responses governments are making and should make to disruptive technology, and how you too can be a futurist.

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>