# AI and You

Transcript

## Guest: David Wood

## Episode 21

First Aired: Monday, November 9, 2020

Welcome to episode 21. In this episode we conclude the interview with David Wood. He is 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*.

In the first part of the interview, we talked about futurism in general and changes in our attitudes towards and perceptions of the future. David introduced himself as a singularitarian and I explained what that meant, and he talked about how groups like the London Futurists work to go beyond superficial gee-whiz views of the future to help us prepare for it in practical ways. Go back in time a few decades and such a group and its activities might have looked more like nebbish sci-fi fandom mumbling in an echo chamber, but now it seems to be exactly the sort of prescient, proactive grass roots research that we need.

Now on to the concluding part of the interview with David Wood.

You're making me think about OpenAI here, the company or conglomeration that said, "We're going to develop AI transparently so that everyone can see the source code. And we will use the "Many eyes make all bugs shallow" principle to develop that to be ethical and less dangerous. And then one of the first things they come up with is GPT-2 and they say, "Oh, this is so dangerous. We're not going to show you the source code.

There are some real issues there. There are some things that probably shouldn't be publicized so widely, such as how to modify the polio or malaria or smallpox viruses to make them more deadly. People are looking, of course, now at the coronaviruses and there is some research on what might happen if there are more deadly. Now, we probably don't want to publicize the details of that. So there are some restrictions. The other issue with OpenAI, of course, is that some parts of it got changed into a more commercially driven way. And you can understand why because it turns out they needed more funding than they had initially expected, and some of the setup of that corporation changed. I'm still hopeful, very hopeful about what OpenAI may do in the world, but it is a warning sign that the mechanisms of capitalism with all their strengths, all the blessings that come out of it also can lead to the change of purpose of individual product lines, to make them more focused on things that are good for minority and good for short-term profits rather than the longer term. So we need to keep on emphasizing the dangers of just pursuing for-profit enterprises, the risks of market failures alongside the risks of political failures, and alongside indeed, the good things that come out of the free market. We need a more

sophisticated understanding, which highlights possibilities for intelligent democratic oversight of both what politicians do, and of what to corporations and groups of corporations do.

Some of what you were saying there reminded me of some of the simplistic associations I was used to making. I'm not saying what you were making very simplistic associations, I'm talking about myself here and that I would think futurist and associate to utopia. Utopia is usually succeeded by the word socialism. And then we're looking at different political systems and then, it seems so much of these visions of the future hinge upon us developing something different to the dog eat dog Law of the Jungle kind of capitalism that we have right now. You seem to be hinting at that how much of your work looks at how we might have to change those systems.

So anybody who's worried about the machine of AI running out of control, doing things that humans don't like, also need to take into account we have potentially the machine of the large corporations already doing that. Large corporations are already superintelligent organisms, which have their own sets of objectives to which other things are often subordinated and so there are risks. Dwight Eisenhower warned about it as long ago as his farewell speech in 1961 on leaving the American presidency. He talked about the risks of a military-industrial complex. And he was not against the military. He was not against the work they were doing. He was not against the set of defense institutes. He said, “All that's important, but it needs to be under political oversight, otherwise, it can become self-perpetuating.” That is the big oil complex. And there is the big tech complex, there is possibly dysfunction in the pharmaceutical industry too, that they are prioritizing treatments which can make them lots of money. As one of their CEOs infamously said, “We did not develop that drug for poor Indians. We developed it for rich Westerners.” So there are risks in each large corporation but this isn't a new insight. I would say in fact, it is a fundamental insight that's been around for 150 years or more. And I would say that possibly the best invention of the last hundred and 50 years, is the idea of the mixed economy that allows lots of leeway to companies when they are entrepreneurial and when they compete, when they seek to find better ways to serve the customers, but then it intervenes when the very large companies start acting as monopolies. So in the past, in American history, there were presidents such as Teddy Roosevelt and Woodrow Wilson, who imposed various breakups of large trusts. And in the more recent history, the companies such as Microsoft and IBM were also subject to some supervision. Microsoft famously said, “We've got to bundle our web browser so closely to the operating system. There's no way one can function without the others.” And the Department of Justice said, “I don't believe you. You've got to make them separate.” And thank goodness they were separate because there's been much more innovation in web browsers from the competition from groups such as-- Well, there’s Opera, there's Firefox, and there's Chrome and of course, Microsoft's own web browser has come along by leaps and bounds. So there are many examples of how society has progressed by this combination. It's not just an invisible hand with things happening in an uncoordinated way, there is from time to time a visible hand that intervenes, and I think we need more of that. I think there is actually a dangerous rewriting of history. Some people have talked about it as amnesia. There's a famous book by a couple of political scientists called American Amnesia, where they say that there's been a deliberate move to try and cover up and obscure what happened in history, and always say that political intervention did bad and it's

always been the free market that's done well. Even if we look at something like health care. We are often proud of the increases in lifespan. The average lifespan has moved up from something like in the 30s to in the 70s. And we can point to biotechnological reasons for that with things like vaccinations and antibiotics, things like heart stents, and so forth. But it has also been due to political interventions that have insisted on hygiene that prevented the use of dirty waters, insisted that milk was pasteurized. There's a retelling of history which I believe in, which is that pasteurization of milk probably saved more lives than many of these other interventions. So that would not have happened by itself. It needed political intervention. It needed coordination and politics is simply to recognize that some goods are public goods that will not be adopted simply by every corporation thinking how they can do best for themselves. They need some alignment. So my work as a future is to try and explain that the route to a better future necessarily involves not just better technology, but also better politics and better economics. And that's part of what will allow us, humans, to flourish much more than now.

And it goes this is coming to the fore right now as people realize that the biosphere and our climate are common goods that have been abused and - a whole other conversation there. I think that just to wind up this theme here for a moment, that you would have [got] general agreement. Most people would say yes, it's a good idea to share the dividend of technology as it increases productivity, as it creates new wealth with everyone, you wouldn't want to withhold it from anyone. But the moment you try saying, “Okay, and how will we do that?”, there are crickets. And because people see, “Well, that must mean you're going to take it from me or who I might aspire to be.” And the best ideas are not well received or not well developed, like Bill Gates’ idea of a robot tax doesn't seem practical. The best idea that seems to be batted around at the moment is universal basic income but we're sitting on the fence about that. So there's not a problem of wealth generation, there’s a problem of wealth distribution.

So I agree with you. This is a difficult topic, and we need to break it down into more manageable chunks. And this is the agile spirit, again, how do you deal with a huge uncertainty? You got to find ways that you can bite off bits, and then string them together again. So part of it is about what is the structure of a better society now, in general terms, and I believe we have to move on from the Protestant work ethic, the Calvinist spirit that if you're not working, you are a second class or a third class citizen; don't deserve anything. We have to embrace the idea of Arthur C. Clarke when talking shortly after 2001 came out, he was a bit alarmed that the computers and AIs in 2001 were seen as malevolent, and he wanted to correct the image a bit. And he said he looks forward to a time when there is total unemployment, rather than total employment, where each of us has much more time to play, was his word, so we can talk about being creative and exploring. So we'd all still have tasks, we would all still be highly engaged, but we'd be engaged in things that weren’t bringing us wealth particularly, because we wouldn't need to live in wealth. Why? Because again, of the robots. The robots will be developing all that we need in this future scenario. They will be doing agriculture, they will be doing a lot of the mining, they will be doing a lot of the manufacturing, they will be doing a lot of the services. And they'll be doing it at a much lower cost and with a much higher quality. So we will not need a big basic income.

We will be able to get by in a very small income because almost everything that we need for a good basic quality of life will be very low cost. So rather than thinking, “How can we redistribute wealth to put much more money in everyone's pocket?” I think the target is “How can we reduce the costs of everything that is important?” And when people point out that robots could in the future, do tasks better and less cost than humans, currently, they set off alarm bells and says, “Well, we'd be out of a job. This would be terrible.” And we have to switch that to “You would be out of a job. That's wonderful.” Now, that's not the end of the conversation by any means but I push this angle as more important even than discussions of universal basic income. I do think that everybody should have some income, but I will make the case that we will need less income than people currently expect once we get [these] technologies of abundance in place. It includes education being more or less free for everyone. Everything we learn currently by paying lots of money to universities or colleges can be accessed online and via communities and via virtual reality headsets, and so forth. It includes much lower-cost health care in which there is preventative health, which prevents more people getting these terrible diseases of aging and chronic diseases at all. Instead, they'll be headed off. A stitch in time saves nine. So that's part of the big vision, then we have to figure out how do we get there and that's another huge topic, but I think we need to clarify the positive vision as an important part of the prelude to this discussion.

Right. And I keep thinking that we may be in a stable situation now where the socio-economics of our capitalist system perpetuates staying in that. It wants to keep returning to an equilibrium and the kind of alternatives that we've been describing could equally well be stable and self-sustaining their own ways but the journey from one to the other involves passing through a very volatile and dangerous territory that requires a lot of courage. Now, I want to go from this to something you mentioned in one of your presentations that seems related where you said, “Disruption requires building an ecosystem with feedback cycles.” And I wonder if you could describe that principle there.

So things that come out of the blue, things that are different ways of doing business, they often don't arrive completely fully formed. It does require a group of different actors, different players to work together. So with mobile phones, you needed the phones, but you also needed the networks. And in the old days, the people who were creating the wireless networks used to joke that GSM stood for God Send Mobile files. GSM was the sort of 2g we now call it. And then with smartphones, you needed not just the phones and the networks, you needed the applications, you needed the developer tools, you needed the verification systems to check that applications weren’t doing bad things, you needed application stores, you needed a lot to come into place. It's the same I think, with many other potential breakthroughs. You need to have several things coming together in order for change to succeed. Even if I look again, at Apple's case. Apple changed the game with smartphones. There were smartphones before, as I know well. My software that my teams helped to write - Symbian OS - were in many smartphones from companies such as Nokia, some in Motorola, some in Sony Ericsson and some in Samsung. Apple changed the game in many ways by having devices which were much more expensive than anyone had previously expected. They were larger screens than people had expected. The previous game was often to shrink the devices. And they focused less on voice, they focused

more on web browsing. But that could only happen because Apple in turn combined several different things together. They combined their previous work on iPods. Some people may remember iPods, the music devices with their focus on ease of use. They also took advantage of the iTunes music distribution system, which turned into an iPhone management system that took advantage of a new kind of touch interface. Previous smartphones had a different kind of touch in the hardware was called resistive touch. And you couldn't easily do the drag and drop and the two-finger tweaks and so on. But Apple put that capacitive touch in their system, and they copied the App Store idea, some of them from Palm and some of them from Nokia but they did them much better. And yes, they didn't have the world's best voice technology, but they could pick it up more easily because voice had become a commodity. So a whole bunch of different things came together to make that possible. So if we look ahead to possible breakthroughs, it's not enough to say, "Well, there's a technological possibility here. One day we're going to get to AGI. One day, we're going to get to Eric Drexler's vision of molecular scale nanofactories. One day we're going to get to amazing new green technologies." You're going to need the lots of things to fall into place. It acquires an architectural point of view, and it requires the kind of open APIs that we spoke about in the past, and so it won't happen automatically. People who complain that Eric Drexler's vision of nanofactories hasn't happened and it's all just due to bad actors possibly subverting some of the funding that was made available to nanotech taking it into different materials instead, which are nanotechnology of a sort, but a different kind of nanotechnology, they have reasonable complaints. But to actually get all the way to that bigger, more exciting, more radical vision, it takes a lot more planning. And I don't think that planning will always be successful simply by venture capitalists investing. Sometimes you need government foresight as well making that structure possible. And if the Western governments don't do it well, I think the leadership of various countries in the East, not just China, but also Korea. Korea has an amazing track record of industrial policy, where going all the way back to the Premiership of [inaudible 44:59] in the late 1960s and '70s. They set very interesting targets for the businesses in Korea to work with. And that's caused companies such as Hyundai, Samsung, and LG, to become some of the world's giants in technology, not just by hardworking individual engineers and entrepreneurs, but by the right kind of political support. So to get to the true potential of technologies, such as artificial general intelligence or for nanotech or for rejuvenation biotechnology, it's going to require architectural vision too, which may not happen inside individual companies. It may require people working alongside visionary politicians.

And what a great segue here to talk about the role of government. The first person I interviewed on this show was Audrey Tang, the information minister of Taiwan. And she gave some amazing information about the Taiwanese policies for dealing with artificial intelligence and the way in which that can support their response to things like the pandemic. And no surprise that she would develop those sorts of things because she is a genius in every sense of the word, and also a highly accomplished computer software developer. Now, you and I met in the British government, the House of Lords in the hearings that the All-Party Parliamentary Group [was] giving with respect to artificial intelligence. Now, I traveled all the way there because I thought it was important to respond to their degree of commitment to understanding artificial intelligence. I'm not aware of many other governments putting in that amount of work

to understand the breadth and depth of artificial intelligence. And yet, I don't know what is going to happen to that work. I don't know how that will be connected to policy or if it's just going to end up in the documents somewhere. Corporations that develop artificial intelligence and other technologies are interested in their own future, their own welfare, their stockholders' interests. They are not necessarily going to try and, as you say, think about the wider impacts of second and third-order effects of what they develop. But if anyone's job is to do that it is a government's. How are we doing in that respect? Are governments, do you think, aware of the responsibility that they have to try in this area?

So there are lots of encouraging signs. I'm a big fan, like you, of the work of Audrey Tang in Taiwan. She does talk about the need to move away from anti-social social media to pro-social social media and some of the systems they've developed there. They're not just neat technologically, they've had a lot of thinking about the implications of what makes people collaborate, as opposed to what makes people get into fierce confrontation online, so that's encouraging. There are people who work with the Obama regime, and towards the end of it, Obama's time in office, he had a number of people who wrote quite visionary things about the ways in which technology would transform. But with the current regime in the US, we've had people like Steven Mnuchin, one of the picks of Donald Trump who is still in place after all these years, who said quite early on, that it was not in his planning at all to consider robots or AIs taking many jobs from humans, and they downplay that possibility. But there are lots of people in America, and people who actually made a movie recently. I forget exactly what it's called. It's something like The Social Network. It's on Netflix. It has people like Tristan Harris from the Center for Humane Technology, and I may again have misremembered that. Hopefully, the show notes will get this correct. It has a bunch of people, some of whom are connected with government who are sharing very insightfully, the risks for what the social media technology can do if it's wrong, and how to restructure it. So there are lots of people around who are transcending the profit motive, shall we say, and it includes some fines and supporters at quite high levels inside the major companies. So Brad Smith, the President, I think Microsoft has said some very encouraging and thoughtful things on exactly this topic. We look back to what's happening in the UK with the All-Party Parliamentary Group, I think they've done a good job at a certain level. They have focused on the need to adapt to significant changes in employment because of technology. So they have said - talking to companies - "You need to up your ability to retrain your employees. The employees will need to retrain more often than in the past. They will need to become more aware of the ups and downs, strengths and weaknesses of AI." So that I like, and that is a good insight. But actually, I don't think they have gone far enough, and they have stuck on the level of "All we need to do is encourage people to retrain and things will broadly be fine", or "Oh, yes, they are focusing on other failure modes inside AI such as the drawbacks of unexpected bias in algorithms, lack of transparency, not using sufficiently representative data." And they're highlighting these things and doing a good job. But I think the bigger risk is from the next wave of AI, which will not just put some people out of jobs and require some retraining, it will put nearly everybody out of a job and ways in which no amount of retraining will allow people to keep up with faster progress of AI. And that is going to require a larger change in society than simply saying, "Let's think a bit more about bias and algorithms."

So what they've done so far is good, and it will go to a certain extent, but they have been a bit shy and nervous of the larger change and I understand why. It's because many of the people who do talk about singularity, talk about the growth of artificial general intelligence, sometimes give the impression that this is inevitable, sometimes give the impression that it's all going to happen by clockwork, and they seem to be insufficiently aware of the social and political dimensions to technological adoption. So at that meeting where you and I both were, I think one speaker who was quite dismissive of the vision of Ray Kurzweil, actually had some fair points that he was making. And although I admire Ray Kurzweil for raising many people's awareness of the possible rapid changes in technology, his views, as outlined in various books, do need to be set into a wider context of the more difficult social and political discussions. And that's what I know both you and I, along with others and Kurzweil himself to now, although he's in sort of a tradition, which sort of typecast him. So the bigger problems cannot be solved at the level of individual governments, even at the UK. It's going to require a larger amount of cooperation, which is why I am distressed at the exceptionalist or isolationist or pro-sovereignty views expressed by many in the British politics. When they say we don't want to be tied down with the slow cumbersome processes of the EU, the EU is thinking about responsible technologies, "Hey, we in Britain, we're going to bound free from that. We're going to be Singapore on the Thames", although actually, if you look at what Singapore has done, it's far from just letting entrepreneurs do what they like. There is an important amount of architectural planning there too. So I think we're going to need to build better links and I do fervently hope that at least some in the UK Government will want to reach out and strengthen links with people in Taiwan and China and America, throughout the EU, throughout the Middle East too. There are groups all thinking about these issues, and they cannot solve these issues in a local way. They need to be addressed [in] a global way. And that's frightening, but we've done it in the past. We have solved other problems globally in the past like preventing too much proliferation of nuclear weapons, chemical weapons, and biological weapons. It hasn't been entirely successful, but we are able to make these international agreements work with sufficient effort. And given the stakes are even higher with AI, then it's even more of an imperative. So I intend to keep working with friends and colleagues who are plugged into that UK's All-Party Parliamentary Group to insist that we don't just think about the next wave of disruption from AI, but we think about the waves of disruption after that.

And I think one of the responsibilities of government is in shaping public conversation and public perceptions. And some of those public perceptions of artificial intelligence right now driven by the Terminator, for instance, because every interview with someone like you or me that the media does seems to be punctuated with a picture of that is driven by a fear of the unknown and a fear that this AI will become superintelligent but have no feeling, have no caring built into it, that it will be a computer that doesn't understand what's important to us. And there's a lot of discussion in work like Roman Yampolsky was on the show and it's predicated on this assumption that AI will not feel. And we don't talk much about the feeling of computers. I think it's an awkward or embarrassing topic to discuss. But at some point in the future, if we don't have computers that can feel but they are super intelligent, then these fears might well be grounded. And this is where I remember that I discovered relatively recently that sentience

is actually not about consciousness or thinking, its feelings. So when we talk about sentient beings, we're talking about feeling ones. What is the prospect for artificial intelligence being developed to have innate ethics, morals, and a sense of caring?

That's the $64 trillion question, I guess. That's a question that you and I could easily discuss for many more than 64 minutes. You can break it down again into lots of different chunks. There are people working on affective computing in which computers recognize human emotions. Computers can manifest the appearance of emotions with smiley faces, and so forth. And they can influence and alter humans' emotions, and they will do that by understanding emotions in a logical way. And that doesn't mean of course, that they actually do have real sentience. Although there are some theories that say that as they become more intelligent, some other sentience will just emerge. We don't know. That's one possibility. Then there are other people who say, "No, you've got to design sentience in. You've got to actually study what is it about the natural brains that generate sentience, and let's copy some of these features and put it into our AIs." And now we want to say why? Why is it so important to have that sentience? And some people say, "Well, if there is sentience, there will be more of a compassion, more of an ethics." But each of these claims, I think, is an open claim. It's not obvious that beings who have more capacity to feel are necessarily more ethical. And even if creatures do have a sense of ethics, it's not obvious that their sense of ethics will dovetail with what we would want to happen - a question of ethical alignment. So there are many humans who don't seem to have much feelings. They are sometimes called psychopathic, or there are other psychological terms, and they don't object to others' suffering. There are others who are fully aware that humans suffer, but they think, "Well, let's do it anyway." It's for the good of the planet, or it's for the good of their own tribe. So merely having feelings, I don't think guarantees that the AIs will do things that are aligned with long-term human beneficial purposes. Merely having intelligence doesn't guarantee it either. We look at some of the nastiest people in the world. I won't mention any names, but you can all pick your favorite nasty people. Some of them are really intelligent, whether they're working in politics or whether they're working in crime syndicates. They are very effective people, they are very intelligent, and they have lots of power. But that power is not used for good, it's instead, as in the famous saying, by Lord Acton, the 19th-century British historian and politician, "Power tends to corrupt; absolute power corrupts absolutely." And so we may have AIs that have sentience, as well as intelligence, but it doesn't necessarily mean that they are going to be motivated to look after humans. They may also understand full well that we humans are distressed at what they are doing, but they may be following along what is programmed into them as some kind of prime directive, intentionally or unintentionally. And their ability to anticipate how their prime directive may malfunction is, of course, a key question. And people like Roman Yampolsky that you mentioned, have given a great deal of thought to this and essentially, I share his conclusion that there is no simple process at all to guarantee that superintelligence will necessarily carry out actions that will benefit the bulk of humanity. And people say, "Well, of course, AI systems become clever. Of course, they're going to want to treat humans well." But I see no such a compulsion. AIs may set themselves entirely different goals. Max Tegmark writes about this too, of course. Max Tegmark from the Future of Life Institute. He imagines if we had programmed AIs with an objective, "Ensure every human gets to heaven,

and you do this by teaching humans to pray and believe in whatever set of catechisms are necessary", the AIs might decide for themselves, "That was an illogical objective. There is no such thing as an afterlife in heaven. What these humans have asked us to do does not compute." It may be that we humans, we try and set goals for the AIs that in a similar way, the super AIs of the future will decide do not compute, in which case all bets are off as to what they actually do.

Wow. As you say, we could go on for a lot more than 64 minutes, and I really want to, but at some point, we have to cut this one off. I want to conclude here [and] ask you: Someone comes up to you and says, "David, I've read your works. I've seen your presentations. I'm convinced you have convinced me that the future needs everyone's attention. I am on board. What should I, whoever I am, Average Joe in the street, what should I be doing about this?

Well, part of it is spreading the word. Whatever has convinced them they need to try and use to convince other people. There's a simple mantra, "We must influence the key influencers. We must educate the educators. We must help all those who in turn can help others to understand the same principles that have led us to realize this is important. And for each person, there will be distractions and inertia and existing commitments and family ties, personal practices that might prevent them from acting on this or seeing it, but let's find clever ways to communicate in multiple ways to help people change their views." So it could be by podcasts, it could be by videos, it could be by cartoons, it could be by pop songs, it could be by poems, it can be by works of fiction, and it could [be] just by good old-fashioned discourse. So let's find different ways to elevate a better understanding which includes improving central sources of data, such as Wikipedia, and other Wikis to have the reliable information about these issues so that when people get confused and unsure, there are systems they can look to which are more reliable and more trustworthy in highlighting what are the key scenarios, the risks of the scenarios, the opportunities or the scenarios. So let's build this bigger community by sharing and developing a richer collective intelligence that will have some Ais in it, but will certainly have lots of humans supporting and nurturing and encouraging each other.

Wow. Well, David, thank you for your particularly eloquent explanation there and that call to action. So how can people listening to you here, find out more about you, what you're doing and help get on board with some of that?

So look up London Futurists, which in this new normal exists entirely online. We no longer have meetings, folks just in the universities in London. So from time to time, there are activities there and also there is a newsletter which I send around with my own views as to some of the most important trends. And if you like you can follow me on Twitter `@dw2`, and from time to time I will share news about new books and new projects which I'm working on.

Terrific. David Wood. Thank you very much for coming on the show.

It's been a real pleasure; you asked great questions.

That's the end of the interview. We sure covered a lot of ground there. A conversation about OpenAI alone, for instance, could easily take several episodes. As an Open Source advocate and community member from way back, I believe that open source is the way to go to develop safe AI; but open source

as a methodology is optimized for several factors, including development speed and maintainability. It is certainly one of the best ways, if not the best way, of developing code that is safe, but is it safe enough when it comes to artificial general intelligence?

The movie David was referring to on Netflix, by the way, was *The Social Dilemma*, which unpacks the effects of social media on individuals and society. I had multiple people pestering me to watch it. It's good for understanding some of the important dynamics of the psychology of social media, but it does dramatize it because, well, computer programmers arguing about algorithms isn't very captivating. That dramatization has consequences, though, particularly when it displays these three guys who are inside the social media platform and are making decisions about how to get a user hooked on their product. While that accurately displays what the algorithms do and how they affect the users, it anthopomorphizes the algorithms to an extent that may have people believing that this was the way the people who designed the algorithms were thinking. But they're not that evil, and they weren't thinking that way. I'm not minimizing the end result, it's just that it was the product of a system, one component of which was programmers and another of which was the algorithm they created. In the same way that we talked a few episodes about Searle's Chinese Room and the counterargument that the *system* of the room was actually doing the thinking, it's that *system* of the programmers, their algorithms, the market pressures, our societal imperatives and the psychology and behavior of the users that conspire together to act as though there were this homunculus at the center of it all that's hell-bent on enslaving the customers. This is just as serious as if there really were a maniacal dictator driving it all, only it's much harder to address because with a dictator you have an assassin or a revolution and problem solved, whereas a system is much more complicated to figure out.

In a little administrative note, our quality control department just discovered that our sound engineering department messed up on the ending to episode 13, the second half of the interview with Paolo Pirjanian. There was about 5 minutes of silence between the end of the interview and the closing remarks, which probably caused everyone to think that the episode had just ended right there. I've fixed it, so our tests show that if you go back and get that episode again, the gap will be gone.

In today's AI headlines, autonomous repair robots will use AI to identify and fix potholes in UK roads. The electric, self-driving robots – which look like a cross between a Zamboni and a tank - are being built by a spin-out company from the University of Liverpool called Robotiz3d and can find small cracks in the road and cover them with asphalt. No tea breaks required.

And in other news, Microsoft has released a public preview of a free app that helps people train machine learning models without writing any code. The Lobe desktop app for Windows and Mac currently only supports image classification, but Microsoft plans to expand it to other models and data types in the future.

I am very excited about next week's episode, because I will be interviewing Pamela McCorduck. She is the historian of artificial intelligence. She realized many years ago that the founding figures of AI had stories to tell and stories about them to be told and that no one had captured those, so she went and interviewed them and gathered this information and put it in a landmark book, called *Machines Who Think*, and a recent sequel, *This Could Be Important*. Those books are absolute delights to read because Pamela is also a novelist and she brings an impish humor to a field that wasn't expecting that. That's on next week's episode of *AI and You*.

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

**http://aiandyou.net**