

Ask A Biologist Vol 092 (Guest Max Tegmark)

Life 3.0

Life in the year 2050, what will it look like? Will it be the age of intelligent computers? Could this happen even sooner? Physicist and author Max Tegmark talks about the future of artificial intelligence and how it might impact life, biological life on this planet. Dr. Biology learns about the future of smart computers. Will they become too smart for humankind? What precautions should we put in place to secure our place on Earth and our future?

Transcript

Dr. Biology: This is "Ask A Biologist," a program about the living world and I'm Dr. Biology. My guest today is Max Tegmark. He's a professor at the Massachusetts Institute of Technology and the Scientific Director of the Foundational Questions Institute.

He's visiting Arizona State University to give the Annual Beyond lecture. He also has a new book called "Life 3.0 -- Being Human in the Age of Artificial Intelligence," which by the way has been on the "New York Times" Best Seller list.

It explores the world of AI, or artificial intelligence, its benefits and potential dark side, while asking the big question, "What sort of future do you want?" Welcome to the show, Max Tegmark, and thank you for visiting with me today.

Max Tegmark: My pleasure.

Dr. Biology: This program, as one would guess, is about life. Mostly biological life that we know, and other types of life forms that might exist that we have not discovered or have not discovered us.

Artificial intelligence, AI, is giving us yet another world to explore. Your book is both an exploration of AI and a cautionary tale. I have to say that before opening your book, the title got me thinking. Life 3.0, if you're writing about Life 3.0, what was Life 1.0, and what version of life are we in now?

Max: Life 1.0 is really dumb life, like bacteria that can't learn anything during its lifetime. I call us humans Life 2.0, because we can learn, which in nerdy computer speak could be thought of as the ability to upload new software into our minds.

I was born a native Swedish speaker, but I thought, "Hey, it's a nice idea to install an English-speaking module into my brain," so I did. It's this ability to learn which has enabled our human species to become the dominant one on this planet.

Life 3.0, which can design not just its own software through learning, but also its own hardware, of course doesn't exist yet. Maybe we should call ourselves 2.1 because we can put in pacemakers, artificial knees, and cochlear implants.

Dr. Biology: The other part of that book is talking about artificial intelligence. A lot of people hear AI. My question was, is artificial intelligence artificial, and is it intelligent?

Max: Artificial intelligence is simply defined as intelligence which is not biological, something that we made. Intelligence itself, people love to argue about how it should be defined.

I like this really broad definition of intelligence as simply our ability to accomplish complex goals, because I hate carbon chauvinism, this attitude that you can't be smart unless you're made of carbon atoms or cells or whatever.

Today, we have machines that are already better than us humans at many very narrow tasks, like multiplying numbers fast, memorizing large data bases, soon driving cars and many other things. The broad intelligence that a human child has, where they can get good at anything, given enough time, does not yet exist in machines.

This is called AGI, artificial general intelligence. It's the holy grail of AI research. Most AI researchers today in polls think we're actually going to get there within a matter of decades, which is why I think it's a good idea to start right now thinking about what we can do to make sure this becomes a good thing and not a bad thing.

Dr. Biology: Yes. When I started reading your book, the introduction there, there's a blue pill and a red pill at the beginning. You make a decision which way you want to go. It all depends on whether you think artificial intelligence is here or coming soon.

I can tell you that I took the pill that led me down, actually a very enjoyable story with the Omegas and a very interesting computer, an AI computer.

When we talk about AI, there are two terms that are used, artificial intelligence and machine learning. Are those the same? Are they different? How do they relate?

Max: Machine learning is simply a special kind of AI where the machine isn't just preprogrammed to do smart things like multiply numbers or play chess, whatever, but it can actually learn and improve itself from data. It's machine learning which has driven most of the breakthroughs in AI in recent years.

Dr. Biology: I think I saw an interesting project from Google along which you would probably call machine learning. It was a fun, little character. It was learning how to walk.

Max: Yeah, the amazing thing is that it was able to learn to walk without ever seeing a video of anybody else walking, without even having a concept. It was just given a bunch of numbers that specify the angles of all the joints in this stick figure. It got a reward point every time it managed to move forward a little bit and eventually it figured how to walk.

It was exactly the same kind of machine learning which recently also enabled computers to beat the world's best players in this Asian board game of Go. I thought, "Hey, let's try chess."

AlphaZero not only crushed all the world's best chess players but more interestingly it also crushed the AI researchers who had spent over 30 years handcrafting AI software to play chess. Machine learning is very powerful.

In some way, what we are doing is reverse-engineering some tricks from biology because of course the reason that human children can get smarter than their parents, is because they can learn for themselves.

Dr. Biology: Does the machine-learning machine actually do it through trial and error? Is that what we are seeing or is it more than that?

Max: There is a lot of different algorithms but all fall under the umbrella of deep learning. We still don't fully understand some of the best methods that our own brain is used to learn.

For example, if you train an AI system to tell apart cat pictures from dog pictures, you have to show them lots and lots of examples until they get good at it. Whereas with human child, you show them one cat first time ever and they can already after that point out other cats for you.

It's a very active field but it's specular how fast things are developing. Just a few years ago, people thought it was going to take decades until machines beat us in Go. Now, it's already happened. Think about how recently it was that there were no self-driving cars and now we have self-landing rockets already.

Dr. Biology: Right. The self-driving cars, I had them in my neighborhood, both the Uber and Waymo. They are both test-driving the cars around. You see them all the time. I'm always wondering, I'm always looking to see if there are hands on the wheel [laughs] because they do have people in them just for safety sake while they are testing them out.

Max: I've read that they're going to stop having people in them in Phoenix, now actually. They just got permission, Google Waymo to take up random passengers with completely autonomous vehicles, I think. I was hoping to get to ASU that way this morning but then turns out it wasn't quite rolled out.

Dr. Biology: You have to use the carbon-based form to get here?

Max: Exactly.

Dr. Biology: You talked a little bit about artificial general intelligence. Let's talk a little bit more about that because that's such an important differentiation between what we call AI and general.

I think you have used the example of first iteration of teaching a computer to play chess. I think it was IBM that beat the grand master. That one was very focused and it only did one thing.

I think you also mentioned that it was very labor intensive because the programmers had to program all that in. Now, we are moving into this ability to do general intelligence and we couple it with machine learning. Where are we now and what do we have to look forward to in, say, a couple of years' time?

Max: There is a fascinating controversy in the AI community about how long it's going to take to reach general artificial intelligence. Most researchers think it's going to happen within decades. If it does, I feel it's going to be the biggest change ever for a life on earth because intelligence gives you power.

What makes us the most powerful life form on earth right now is not that we have stronger biceps than tigers or sharper claws than tigers. It is that we are smarter than the tigers. If we build machines that can do everything better than we can, it unleashes this enormous power that can be used either for good or for bad.

I was at the hospital recently and told that a good friend of mine had an incurable cancer. It's of course not incurable according to the laws of physics. That just meant we humans weren't smart enough to figure out the cure.

If we can build our artificial general intelligence and amplify our own intelligence with machine intelligence, it absolutely opens the key to cracking all of the most difficult problems of today and tomorrow and helping humanity flourish like never before.

At the same time, we obviously need to be careful, because if whoever's controlling this power is somebody or something whose goals aren't aligned with ours, that's bad news.

The key thing right now is not to freak out and start quibbling about how worried you should be, but rather to ask, "What are the useful things we can do today to try to make sure we steer this technology in a good direction?"

Dr. Biology: On that topic, you have a well-known supporter, Elon Musk. It was a surprise to me to know that he was such a big supporter, because in the news when I see the headlines, it looks like he's the doomsayer.

"Oh, no, no, we don't want to go down that path." It's a little different in that story, isn't it?

Max: Of course. Journalists know that fear sells. They will always, for click bait, try to twist things into the most doomy, possible way. Fact of the matter is Elon Musk is a very optimistic visionary who likes to think long-term.

He uses AI to land his rockets. He uses AI to fly his car towards Mars and to drive his Teslas. It's exactly because he's thinking more long-term than most other people, that he sees how important it is that we get this right and don't screw it up.

Dr. Biology: Let's continue on that path. There are two questions I'm going to ask, and you can go with the first question and the second one, or you can flip them around.

One has to do with, what are the positive results you see coming with AGI, when we use AGI, and what are some of the concerns you have about AGI?

Max: Basically, everything I love about civilization is the product of intelligence. If you list all of the top problems that we have failed to solve so far, AI can help with all of them. Curing diseases, helping us live longer, healthier lives, new technologies, tackle the climate challenges we have, help eliminate poverty, you name it.

At the same time, intelligence itself shouldn't be confused with something which is morally good. Intelligence is not evil, intelligence is not good either. Intelligence is just the tool -- artificial intelligence -- that lets you accomplish goals.

That means it's not enough to just make something really smart. You also have to make sure that the goals built into the system are such that it will do good things, not bad things.

For example, right now, we're on the cusp of starting an arms race and lethal autonomous weapons, whose goal is very explicitly to kill people in novel ways. Most AI researchers are trying to get an international ban on this, just like most biologists pushed hard and successfully to ban bio weapons. We'll see how that goes.

There's also the business of money. Obviously, AI has the potential to grow our economic pie enormously, with AI producing, and robots producing, even more goods and services. If we can't figure out a way of sharing this great new wealth so that everybody gets better off, then shame on us, I feel.

Many economists that argue that the rise in inequality that we're seeing today, which is causing so much more polarization in our society, is in part driven by this. It used to be the technology eliminated really crappy jobs during the industrial revolution and those people could educate themselves and get better jobs.

Now we're seeing AI more and more eliminating really good jobs, and those who lose them have to switch into more poorly paid jobs, because they're taking away, not muscle power, but brain power.

Dr. Biology: When I was in high school, and we won't talk about how long ago that was, I had a really wonderful English teacher. He had a thought question for us, "In the future, what if there aren't enough jobs to go around? Don't worry about money, but what would you do?"

With AI and machine learning, you can see that there are going to be less and less jobs. Even if they're "the crappy jobs" or even some of the better jobs, how are we going to have AI help humankind in a way that benefits everyone?

Max: Yeah, absolutely. Jobs, of course, don't just provide us with money, they also provide many people with a sense of purpose and meaning in their lives, and with the social network as well.

If in the future, people will live in a jobless society, it's crucial, first of all, that the government figures out a way of bringing in enough tax money and using that so that people actually have the income they need, and also to make sure that we create a society where we can really flourish.

This is a question that's so challenging. We can't just leave it to AI nerds like myself to figure out. We need economists, psychologists, and so many other people to join this discussion about what kind of future we want.

I often get students walking into my office for career advice, and I always ask them the same question, "Where do you want to be in the future?" If all she can say is, "Oh, maybe I'll have cancer, maybe I'll have been murdered," that's a terrible career planning strategy.

I feel we as a society are doing exactly that every time we go to the movies and watch a Hollywood flick about the future, because they're always dystopian. "Blade Runner," you name it.

We really, instead, need to envision positive futures that we're excited about, because that's going to make it much more likely that we're going to get that future.

Dr. Biology: On top of that, I think it'd be far more sustainable. If it's us versus them, that's a lot more of a problem as we go.

Max: Exactly. When people only think about disease, and threats, and fear, it gets very polarized. When people have a shared common goal, whether it's to go to the moon or something else, that brings people together, fosters collaboration and a sustainable, flourishing society.

Dr. Biology: When we were talking about Elon Musk, we did talk about what the press can do with him, but he's gotten you on a path with AI that's an intriguing one. I just was interested in his challenge to you, that basically how we can work with AI.

I think the phrase was, for you, you wanted to be proactive, rather than reactive.

Max: Exactly. I'm optimistic that we can create an inspiring future for life with technology, as long as we win this race between the growing power of the technology and the growing wisdom with which we manage it.

In the past, we've always managed to stay ahead in this race through the strategy of being reactive. We invented fire, screwed up a bunch of times, [laughs] invented the fire extinguisher. Invented the automobile, screwed up a lot, then invented the seat belt, the traffic light and the airbag.

With really powerful technology, like nuclear weapons or super intelligent AI, we don't want to learn from mistakes anymore, that's a ridiculous strategy. It's much better to shift into being proactive. Plan ahead, think through the things that can go wrong so that they will go right.

Dr. Biology: As a biologist, there are certain protocols that we have. When we've been doing some things in the lab, we're very careful about things not getting out into the wild, we call them. An example of that would be Africanized bees.

It was a very worthy research project, but they got out into the wild. They're basically doing their march all the way out through the Americas. With AI, is there a concern about the really smart machine getting out in the wild and actually taking over?

Max: Once we get closer to AGI, of course, there will be. By definition, AGI can do everything as well as we humans can. Yes. Even long before that, there is a possibility that someone developed some very powerful malware for cyber-attacks or whatever, and it gets out.

AI researchers have to switch into the same safety mentality as biologists, where they really have good containment, both to prevent things from leaking out onto the Internet and to prevent other people from hacking into their lab and getting their tools.

We saw, for example, that the US Government was quite sloppy in some of their own hacking tools that they had built, were used now against us, and against companies.

Dr. Biology: There is a cautionary tale. Be careful what you create.

Max: Yeah. [laughs] I still sometimes come across this attitude, that, "Hey, let's just build the machines that can do everything better that we can. What could possibly go wrong? No need to take any precautions or have any discussions about what could happen."

That kind of attitude, that we're just going to try to make ourselves obsolete as fast as possible, is frankly just embarrassingly unambitious and lame. Humanity is traditionally a very ambitious species. We should aspire to much more than that. We should think very carefully about what kind of exciting future we want, and then try to create it.

Dr. Biology: Before any of my guests leave, I have three questions I always ask. We'll launch into those. When did you first know that you wanted to be a physicist? What was the aha moment?

Max: It was when I read this popular book, "Surely You're Joking, Mr. Feynman!" which wasn't about physics at all. It was about picking locks and having crazy adventures.

You could read between the lines that he loved physics, which intrigued me, because physics used to be my most boring subject in high school. I decided to investigate what I had missed about physics.

I started reading "Richard Feynman Lectures of Physics," volume one, and it was the closest that I've ever come to a religious experience. I was like, "Wow! This is so amazing." That's how it started.

Dr. Biology: You must have had some mathematic background, because if you pick up Feynman's lectures, those are not for the faint of heart.

Max: I was always very curious about stuff. I loved math and trying to figure things out. Physics, I had just been taught in school, was a boring cookbook recipe list for how to calculate certain things, which just weren't interesting.

Feynman helped me realize that what physics really is, is the ultimate detective story to try to figure out the deepest mysteries of reality, except you don't just get to read about it, you get to be the detective.

Dr. Biology: Now, I'm going to take it all away from you. I'm going to take away your physics. You can't be a physicist. Typically, my guests that teach at a university will fall back into teaching. I'm going to take it all away, and I'm going to say, what would you be and what would you do?

Max: You're making it too easy for me. I would be an AI researcher, because that's actually what my technical research at MIT has been on for the last few years now. I've been just so fascinated by AI that I've built up a group in our department with physics and math students who are doing AI.

Dr. Biology: I made it easy for you. I'm going to make it harder. I'm going to take that away. No AI, no physics, no teaching. What passion do you have? What have you always wanted to do that maybe you've never done?

I'm empowering you more than you typically would get.

Max: [laughs] By taking everything away.

Dr. Biology: Yeah.

Max: I'm also very passionate about applying the knowledge we get from science to make the world better.

If you took all those things away, I would probably spend much more time on the Future of Life Institute, this non-profit where I work with Elon Musk, to try to make people think more long term, beyond the next election cycle and realize how amazing the potential is for the future life if we use tech wisely instead of just to start an accidental nuclear war or make ourselves obsolete.

After 13.8 billion years, our universe has woken up with this self-aware quirk blobs, like us that can experience emotions, and curiosity, and passion, and so on. It will be such a bummer to squander this when life still is so rare in our cosmos.

We have such incredible potential for life, not just on earth for the next election cycle, but for billions of years and throughout this amazing universe. I would probably, if you took all my day jobs away, spend more time on that.

Dr. Biology: You like writing? This is your second book.

Max: I love writing, yeah.

Dr. Biology: In the book, that path I went down then, the little story as you call it, did you always plan on having the little story? Is that how you introduced the book or did you think about it afterwards?

Max: That story was actually the first thing I wrote in the book. I never write my books starting with page one. This was going to go on chapter four. I wrote chapter four, and then five, and then six.

Later, my editor said, "Hey, why don't you just take that out and open the book with it?"

Dr. Biology: That's why you have editors. I think it was a brilliant idea.

Max: I did it because people are so hung up about robots, but it's not robots that's the big deal. Hinges and actuators and motors is an old technology. It's intelligence, that's the big deal. I wanted people to understand that pure intelligence, even without a robotic body, gives you incredible power.

Dr. Biology: Without a doubt. I could see that going into a movie very quickly. I was like, "Oh my heavens!" [laughs] because you really didn't end it.

Max: That was on purpose, because I did it to get the reader thinking about how it might end, and whether they would want this to happen or not want it to happen. If so, how they would want it to end. Ultimately, this is a story where we as a species get to write our own ending.

Dr. Biology: Good or bad.

Max: Yeah.

Dr. Biology: The last question is, what advice would you have for a future physicist, cosmologist or perhaps someone who has done this as a hobby? Maybe they love the idea of AI. They've been exploring it in the news, hopefully in the better parts of the news, and reading books. What's your advice?

Max: My first suggestion is, look beyond the stuff that they tell you to learn in school, and maybe take some online courses, read some great popular science books to get more of a fascination than understanding of the basics.

Use that to get to go somewhere, to some university or wherever, where you can collaborate with a team of like-minded people, because that's the most fun and easy way to really get into orbit with these things.

Dr. Biology: Max Tegmark, thank you very much for visiting with me.

Max: Thank you so much for having me. It's a pleasure.

Dr. Biology: You've been listening to Ask A Biologist and my guest has been Max Tegmark, Professor at the Massachusetts Institute of Technology and the Scientific Director of the Foundational Questions Institute.

Now, if you haven't picked up a copy of his book, I suggest you put it on your list. It's called Life 3.0 -- Being Human in the Age of Artificial Intelligence. Professor Tegmark is in town to give a lecture for the Beyond Center, and I look forward to the event tonight.

The Ask A Biologist podcast is produced on the campus of Arizona State University, and is recorded in the Grass Roots Studio, housed in the School of Life Sciences, which is an academic unit of the College of Liberal Arts and Sciences.

Remember, even though our program is not broadcast live, you can still send us your questions about biology using our companion website. The address is askabiologist.asu.edu, or you can just Google the words, Ask A Biologist.

I'm Dr. Biology.

Transcription by CastingWords