



Life 3.0: Being Human in the Age of Artificial Intelligence

Max Tegmark

[Download now](#)

[Read Online ➔](#)

Life 3.0: Being Human in the Age of Artificial Intelligence

Max Tegmark

Life 3.0: Being Human in the Age of Artificial Intelligence Max Tegmark

How will Artificial Intelligence affect crime, war, justice, jobs, society and our very sense of being human? The rise of AI has the potential to transform our future more than any other technology--and there's nobody better qualified or situated to explore that future than Max Tegmark, an MIT professor who's helped mainstream research on how to keep AI beneficial.

How can we grow our prosperity through automation without leaving people lacking income or purpose? What career advice should we give today's kids? How can we make future AI systems more robust, so that they do what we want without crashing, malfunctioning or getting hacked? Should we fear an arms race in lethal autonomous weapons? Will machines eventually outsmart us at all tasks, replacing humans on the job market and perhaps altogether? Will AI help life flourish like never before or give us more power than we can handle?

What sort of future do *you* want? This book empowers you to join what may be the most important conversation of our time. It doesn't shy away from the full range of viewpoints or from the most controversial issues--from superintelligence to meaning, consciousness and the ultimate physical limits on life in the cosmos.

Life 3.0: Being Human in the Age of Artificial Intelligence Details

Date : Published August 29th 2017 by Knopf Publishing Group

ISBN : 9781101946596

Author : Max Tegmark

Format : Hardcover 364 pages

Genre : Science, Nonfiction, Technology, Philosophy, Artificial Intelligence



[Download Life 3.0: Being Human in the Age of Artificial Intelligence pdf](#)



[Read Online Life 3.0: Being Human in the Age of Artificial Intelligence pdf](#)

Download and Read Free Online Life 3.0: Being Human in the Age of Artificial Intelligence Max Tegmark

From Reader Review Life 3.0: Being Human in the Age of Artificial Intelligence for online ebook

Wendy says

“Life 3.0: Being Human in the Age of Artificial Intelligence” which I won through Goodreads Giveaways is a fascinating subject that seems like the stuff of a science-fiction novel. It begins with an imaginative “what-if” premise with the tale of the Omega Team who, with a strong commitment to helping humanity secretly build an AI called Prometheus. With security measures in place, this ultra-intelligent machine not only makes millions for its parent organization but transforms the world positively and negatively.

With the creation of Artificial Intelligence closer than we can imagine, Max Tegman enables a layperson like myself to look its possibility, form (i.e. machine, cyborg), pros and cons, as well as impact on humanity (i.e. jobs, laws, weapons) with a perspective of its future potential. A uniquely interesting topic especially one which would affect mankind on a global scale, he deals with it honestly pointing to the most controversial issues near the end of the book.

“Life 3.0: Being Human in the Age of Artificial Intelligence” is a book that I intend to read again as Max Tegman’s discussion of Artificial Intelligence is very pertinent and a loaded topic not only for our generation but as we look ahead to the future.

Graeme Roberts says

Probably the best book so far on the possibilities and dangers of artificial intelligence. The grown-up boy genius, Max Tegmark, is ebullient, full of energy, and very charming. What might be considered name dropping in another author is readily forgiven, as he pulls us into the excitement of this rapidly developing field, and his cofounding of The Future Life Institute to ensure that safety measures are adopted to stop AI from dispensing with its dumb old human forebears.

The prelude contains a very plausible but imaginative fictional account of the Omega Team and Prometheus, its powerful AI. Storytelling brings the technology and its human implications vividly to life.

Some chapters are intensely exciting, opening new intellectual vistas with the turn of every page. Others, including *Our Cosmic Endowment: The Next Billion Years and Beyond* and *Consciousness* are a touch turgid and tedious, but I never found the journey less than worthwhile.

Lor-El says

Unlike the author, I will try to be concise and make my point clear:

1) There is so much name-dropping (authors, books, theories, Ivy League Universities, tv shows, movies... you name it, it's gonna be there) and basically it seems like a secondary literature review rather than an

original work.

2) Some chapters feels like fillers, put there just to make the book thicker, they add little to no useful information on Humans and AIs whatsoever.

3) The long awaited (like, 300+ pages awaited) chapters on all the ethical questions and consciousness, which would have made the book worth reading for me, are just a scratch on the surface and do not delve into the depth of these multilayered issues. A show like "Person of Interest" (which, although the Prometheus-AI takeover described in the book is uncannily similar to the one perpetrated by Samaritan, is not even once cited in the book) did a better job at dealing with both ethics, machine learning and how human deal with two sentient and very different ASI. And I know... I know, "that's just sci-fi" but ultimately, so is this book.

Manny says

Last week, I read Nick Bostrom's groundbreaking *Superintelligence*, an extremely serious, well thought out look at the dangers of creating real artificial intelligences. It left me feeling more than a little concerned: despite working in AI myself, I had not fully appreciated how scary it is. I've just finished this book, written about three years after Bostrom's, and now I'm even more concerned.

Bostrom's book has the air of being primarily intended for senior policy-makers in industry and government. It's dry and formal in tone, relentlessly footnoted, full of difficult words, and seldom goes out its way to try and entertain you. It appears to be the product of a great deal of diligent work, and the dominant note is one of a dire, Cassandra-like warning. But despite the fact that it quotes Bostrom on every other page, Tegmark's book is pretty much the opposite: it's manic and chatty, dumbed down to one of the lower common denominators, poorly structured, and full of winsome autobiographical revelations and engaging little science-fiction stories. I could have taken all that - I've read Tegmark's previous book, and I'm familiar with his style - but what really worries me is that he's trying to make us feel *optimistic* and *hopeful*. It'll be okay! He's founded his Future of Life Institute, all sorts of prominent AI people have joined up, they're organising conferences on AI safety and an incredible ongoing web conversation which anybody can join, and pretty soon they'll have the answer!

Well, if you believe that you'll believe anything. As Bostrom explains, the problems here are not technical issues that can be addressed by good engineering. They are foundational philosophical questions that have been discussed for thousands of years without a great deal of progress having been made. It's readily conceivable that there *are* no solutions to them, and there's nothing available right now which even looks vaguely like a possible road towards a solution. Anyone expecting a magic fix to turn up over the next couple of decades is probably also expecting it to be delivered in person by the Easter Bunny riding a pink fluffy unicorn. And yet Tegmark goes on and on about his preferred future, where we become spacefaring AIs who head out, à la Olaf Stapledon, to transform the universe into one huge superintelligence which, after several billion years of development, will be able to perform ten to some very big number of computations per second. He spends inordinately many pages explaining various tweaks, all of which will no doubt be invalidated by the next major discovery in physics, that mean that it will be ten to this big number rather than ten to some slightly smaller big number. Why we're supposed to do it is never once explained. Stapledon's Cosmic Mind gets to make contact with the Star Maker, which at least is a worthy goal, but what Tegmark is hoping for is beyond me. I guess it's just, you know, cool. But all the same, we're supposed to stake the whole future of humanity on this bizarre and geeky dream.

If Tegmark were a moron, I would be slightly reassured. I figured out some time ago that many books are written by morons. But he's a very clever person, who, at least on his own account, has considerable influence over the Secret International Network of Very Clever People. I suppose he may be fronting the superintelligent AI who's currently taking over the world; I don't know if that would make it better or worse, but at least it would make it more comprehensible. Aaaargh!

To conclude, I can't do better than to quote a couple of passages which frequently occurred to me while reading *Life 3.0*. The first, taken from the final section of Bostrom's book, is one conspicuously not cited by Tegmark:

Before the prospect of an intelligence explosion, we humans are like small children playing with a bomb. Such is the mismatch between the power of our plaything and the immaturity of our conduct. Superintelligence is a challenge for which we are not ready now and will not be ready for a long time. We have little idea when the detonation will occur, though if we hold the device to our ear we can hear a faint ticking sound.

For a child with an undetonated bomb in its hands, a sensible thing to do would be to put it down gently, quickly back out of the room, and contact the nearest adult. Yet what we have here is not one child but many, each with access to an independent trigger mechanism. The chances that we will *all* find the sense to put down the dangerous stuff seem almost negligible. Some little idiot is bound to press the ignite button just to see what happens.

The second passage, which has been one of my favorites since I was a teen, is from another exceptionally clear-sighted book, Kurt Vonnegut's *Cat's Cradle*:

And I remembered *The Fourteenth Book of Bokonon*, which I had read in its entirety the night before. *The Fourteenth Book* is entitled 'What Can a Thoughtful Man Hope for Mankind on Earth, Given the Experiences of the Past Million Years?'

It doesn't take long to read *The Fourteenth Book*. It consists of one word and a period.

This is it:

'Nothing.'

[And on further consideration...]

As you can see, I found *Life 3.0* extremely annoying. But at the same time, I also appreciate the point Tegmark makes, that it doesn't help just to say that we're all going to die. Some more positive approach is required.

It seems to me, given the analysis in Bostrom's very sensible book, that there is little chance of making superintelligence 'safe' in any normal sense of the word. If we create a superintelligence, it will almost certainly replace us before long. So it follows that we have two realistic choices:

- a) Do not create a superintelligence.
- b) Create a superintelligence that we could happily envisage replacing us.

I think most people will go for (a). If we are even going to consider (b), I would say that the priority should be to develop AIs with qualities that people will value, and which they feel are characteristic of what it truly means to be human. Some obvious candidates include creativity, humour, curiosity, honesty, appreciation of beauty, empathy and love.

Tegmark briefly mentions "creativity", though his analysis of the already renowned move 37 in the second AlphaGo-Lee Sedol game is as superficial as the rest of the book. I see no real discussion of the other qualities.

Mohamed al-Jamri says

?????? 3.0

???? ??????? ?????? ??? ?????? ???????:

1. ?????? 1.0: ??? ?????? ?? ?????? ?????? ?????? ??? ?????? (?? ????) ?????? ??? ?????? ?????? ?????? ??????
????? ??? ??? ?????? ?????? ?????????? ??? ?????? ??? ??????????
2. ?????? 2.0: ??? ?????? ?? ?????? ?????? ?? ?????? ?????? ?? ?????? ?????? ??? ?? ?????? ?????? ?????? ?????????? ???
????? ??? ?????? ?????? ?????? ??? ?????? ??? ?????? ??? ?????? ??? ?????? ???
3. ?????? 3.0: ??? ?? ?????? ?????? ?????? ?????? ??? ?????? ?? ?????? ?????? ?? ?????? ?????? . ?????? ??? ?????? ?????? ??
?? ?????? ?????? ??? ?????? ?????? ?????? ?????? ?????? ?????? ?????? ?? ?????? ??? ?????? ?????? ?????? ?????? ??????

Tashfin Awal says

I received this book for free through Goodreads Giveaways and have chosen to give my honest opinion about it.

This book was a phenomenal read! It opened my mind up to so many new things, it was truly very insightful! The arguments and ideas presented in this book develop some current well-known facts as well as some new

theories, all of which are incredibly interesting! The lengths to which arguments are presented and the details greatly help develop the book.

Manuel Antão says

If you're into stuff like this, you can read the full review.

Deep Learning Architectures: “Life 3.0 - Being Human in the Age of Artificial Intelligence” by Max Tegmark

“Life 3.0, which can design not only its software but also its hardware. In other words, Life 3.0 is the master of its own destiny, finally fully free from its evolutionary shackles.”

In “Life 3.0 - Being Human in the Age of Artificial Intelligence” by Max Tegmark

See how good your PC is as it ages or you want to install a better graphics card, does the driver play nice with everything? Are you competent enough to sort it out or are you the sort of person who offloads that to IT? The guys in IT are like ducks or swans, all seems serene on the surface but underneath they are paddling hard to stay afloat. They are one badly written security update away from disaster. Do they install the latest security patch or wait for others to see what happens? Also, the more complex a system becomes the more subject it is to critical failures from minor changes, the more they become like having 100 spinning plates on the go at once. If your bank's computer goes belly up just as the proceeds from your house sale are sailing through the system from one solicitor to another is there enough of a data trail to prove it existed? Do you feel lucky? In this day and age, when the state-of-affairs is like the one I'm describing above, can we still talk about AI?

If you're into Computer Science, read on.

ScienceOfSuccess says

tl;dr My animated summary of Life 3.0 is available here: <https://youtube.com/watch?v=9XIcOo5mwQ4>

We are approaching times when machines start to understand our world. There is a possibility, that in the near future we will be working with Artificial Intelligence as equal partners. This idea divided people into two groups. Some people expect androids to be our slaves, and others think that people will be working for AI. How should we approach this new age? Should we destroy all electronics, or work harder to make AI more powerful as soon as possible?

Max Tegmark talks about every possible argument and every point of view regarding AI. It's hard to find the main conclusion, because we have only theories. I found that a lot of those possibilities made me think more about my future and life overall.

The author did an amazing job explaining the most likely outcomes with comparisons so simple, that even grandmas lacking technology knowledge could understand it. I never enjoyed science fiction, but Max used popular stories to better show ways AI development could go and kept it simple, steered clear of exaggerations.

This is a must-read, whether you are a supporter or an opponent of AI.

Radiantflux says

68th book for 2017.

This book is somewhere between a 2-and-3 stars for me.

The book focuses on the long-term dangers of General Artificial Intelligence. The sort of problems that might occur with the great-grandson of HAL3000 in the distant future. There is no discussion of short-to-medium-term dangers from AI destroying millions of jobs. This book is firmly focussed on the danger of the coming age of superminds.

One of the perceived dangers mentioned multiple times is of the form "how do I keep on living when I am no longer the smartest person in room?" the sort of issue I guess most non-tenured MIT professors don't worry about.

The constant name dropping/adulation of various researchers was irritating/distracting (e.g., my wife and I had a delightful dinner with Elon Musk where we discussed etc etc). His description of Tononi in the consciousness chapter was so over the top it had me spitting out my coffee in laughter.

I wouldn't have minded the deep future speculation if it just felt richer. There are no intellectual fireworks here. This is a pretty pedestrian plod through territory that has been better covered in many scifi works. The topic I know best, consciousness, was covered in a superficial and quite biased fashion, which makes me worry that areas I don't know so well were similarly short-changed.

I suspect Superintelligence by Nick Bostrom would offer a far better read.

David Yoon says

Tegmark is an exuberant AI cheerleader awash in the unbridled nerdy enthusiasm of an inevitable post-human future. To his credit the book, a reflection of the work he's doing out in the world, attempts to broaden the discussion around AI to something more than wondering if sentient robots will kill us all.

His prelude on a plausible AI trajectory is compelling and thoughtful stuff and I loved how it expanded the way I think of AI's progress. The exploration into considering whether super intelligent AI become zookeepers or benevolent dictators or enslaved gods is great too. But thinking about the philosophical considerations of consciousness, intelligence and evolution left me cold and the name dropping, back slapping historical narrative added yet another element to this unbalanced read.

In here is a fascinating exploration of what AI could mean for the world, it's just buried under a lot of wonky, wordy stuff that obscured the picture I was trying to form. Maybe I've been spoiled by more narrative, bite-sized, non-fiction - I want the abridged version of this book.

Niklas Laninge says

By far the best book i have read on the topic. To approach life and AI from a physicists perspective really sets this book apart from say Superintelligence, 2nd Machineage and Humans need not apply. Also, the fictional aspects really makes this book a bit of a page-turner.

Gary says

Nothing in this book was original. All of the many topics covered in this book were covered in other books that I have read previously. There was definitely a tedious feel to each of the chapters. I think there is no more important or interesting topic then super AI and the author is right when he wants to highlight the topic, but, please tell me things I don't already know.

The author started each chapter by telling me something that I had already knew, then he would tell me almost nothing more than what I knew, then he would summarize what he had said as if I had not already had read it elsewhere or within the chapter. Overall, the lack of depth of each topic presented made for a tedious presentation.

I'm not even bothered by the things I think he got wrong. He makes a categorical error on the nature of our feelings. He thinks of them as things or objects which give us our subjective reality (subjectivity is his standard for self awareness, consciousness). I don't think our feeling are things. He also appealed to S1 (quick) verse S2 (reflective) memory states. He had brought that up right after mentioning Popper's falsification criteria. I found that somewhat ironic because S1 S2 definitely fall into the pseudo-science realm because there is no experiment in principal that I could design that would refute that S1 S2 paradigmatic hypothesis. I had a lot of these kinds of problems with this book, but the book failed mostly because it was superficial and it did not give a compelling narrative to wow the reader.

There is no more interesting topic than super AI. It gets at the heart of what it means to be human, what our purpose and meaning is, what is consciousness, is there other complex life in the universe (yes, super AI relates to the Fermi paradox with self replicating von Neumann machines), and even whether or not we are living within a simulation (by all means, check out Neil deGrasse Tyson's 'are we living in a simulation' hosted debate. It is well worth the two hours), and there is a way to tie Martin Heidegger into the discussion because his ready-at-hand, present-at-hand, and dasein schema means that solipsism is nonsensical and conscious (self aware) machines are not possible (his opinion, not mine), or one could mention Hegel in detail on these topics because of his relating the in itself to the for itself to the other (or in other words the subjective to the objective to spirit, or self awareness) but this author definitely didn't have much philosophy behind him and it showed by his lack of depth regarding philosophy or philosophy of science.

I have a weakness for books about AI. There should be no way that they bore me, but this author did that by not being able to connect the dots and by only telling me things that were overly familiar and having no real philosophical background as a foundation. (Read Bostrom's , or Kurzweil's, or Pedro Domingos' book, or

either of Harari's books. Each (except for Domingo) was mentioned in this book).

Brian Clegg says

I have to confess that my first reaction to this book was not anything to do with the contents, but trying to work out if there was something really clever about the way the book's title is printed on the spine in white on cream, so it's illegible - would it be, for example, a subtle test of human versus artificial intelligence (AI)? However, that was just a distraction.

Max Tegmark is an interesting and provocative thinker in the physics arena, so I had high hopes for what he'd come up with exploring the future of AI and its relationship to human beings. It's worth explaining that the title of the book refers to three 'levels' of life where 1.0 is 'can survive and replicate' (e.g. bacteria), 2.0 is can design its own software (e.g. us - where 'software' refers to our concepts, ideas and extended abilities such as language) and 3.0 is can design its own hardware, enabling it to transform itself more directly and quickly than our creativity enables us to do.

The book starts with a bit of fiction, which I'm usually nervous about, but it actually works very well, as it's presented more like a non-fiction description of a business development rather than attempting all the quirks of fiction. In it we have a semi-plausible description of how a company that succeeds in producing a self-enhancing AI could take over the world. And this is genuinely thought-provoking.

So, early on, I was convinced I was going to love this book. But unfortunately there is an awful lot of futurology in here (aka guesswork) and like all futurology, Tegmark's can be frustratingly specific about things that we are highly unlikely to be able to predict - though at least he recognises this is the case and points it out. He covers the various ways a super-intelligent AI could develop, whether it would become a rogue, how we'd interact with it... and then plunges on into more and more dramatic speculation, including a chapter that looks forward 'the next billion years and more.' Forgive me for feeling a bit 'So what?' about this.

There is no doubt the whole business of super-AIs is an issue that needs thinking about and discussing - and Tegmark does this in an approachable and engaging fashion. It probably needs reading alongside Nick Bostrom's Superintelligence to get a well-rounded picture, though. It would have helped if it had been significantly shorter - it came across as being long because it was the kind of 'big book' that has to be chunky, and I think it would have been a lot more effective at half its length. One particular section that was ripe for trimming had a long list of scenarios, each of which was then worked through - dull reading, I'm afraid.

In the end futurology is a bit like being told about someone else's dream. It probably seems fascinating to them, but it's hard to get too excited about it as a reader. Life 3.0 is an interesting book, but feels rather like a pet project, rather than a strong popular science title.

Andrei Khrapavitski says

Just finished listening to an audio version of "Life 3.0: Being Human in the Age of Artificial Intelligence," a new book by Max Tegmark. His "My Mathematical Universe" is one of my favorites, so I was really looking forward to his new book. And he didn't disappoint. This is a gripping text for anyone interested in AI and the

future of life on our planet and beyond. Without a doubt, this is the most important conversation of our times. If you fail to see why it is so important, consider this. Most AI experts and industry insiders predict that artificial general intelligence (human-level and above) will become a reality either in a couple of decades or within this century. If someone told you that aliens are heading toward our planet and will land in some 20 years, I guess a lot of people, I mean really a lot of people, would start to freak out. Politicians, the academe, etc. would start thinking hard what to do about it. We are in a similar situation with AI research. Considering the incredible progress made in the field, it is time to think how to make sure AGI, if/when created, is beneficial to humanity.

If you are wondering about the title of the book, Tegmark posits that Life 1.0 is life where both the hardware and software are evolved rather than designed (for example, bacteria). Humans, on the other hand, are examples of “Life 2.0”: life whose hardware is evolved, but whose software is largely designed. By software Tegmark means all the algorithms and knowledge that we use to process the information from our senses and decide what to do—everything from the ability to recognize your friends when you see them to your ability to walk, read, write, calculate, sing and tell jokes.

Finally, Life 3.0 designs and upgrades both its software and hardware.

In summary he divides the development of life into three stages, distinguished by life’s ability to design itself:

- Life 1.0 (biological stage): evolves its hardware and software
- Life 2.0 (cultural stage): evolves its hardware, designs much of its software
- Life 3.0 (technological stage): designs its hardware and software.

In fact, considering the current progress with prosthetics, cochlear devices, etc., he says, we are more like Life 2.1 already. But he thinks it is more likely that AGI will be created faster than any cyborg-style mind uploading becomes a reality (for description of that vision, see an interesting book "The Age of Em" by by Robin Hanson).

Tegmark begins his book with a pretty realistic scenario of a superintelligence break-out. Reads like a sci-fi thriller but is in fact much better than what Hollywood has come up with so far on this subject.

He then explores the current state of research into machine learning and some breakthroughs in the field. Then he tries to imagine the near and more distant future. It takes a physicist to write a compelling vision how far life can progress if limited only by the laws of physics. This part of the book is truly mind-boggling even if most of it can hardly be achievable due to various limitations and possible cosmocalypses (also described by the author).

Then he explores the subject of consciousness. Many people view AGI as our descendants. Even if they choose to eliminate us, they will live on and continue the story of life in our part of the observable Universe. Well, what if what we create are zombies without any consciousness? Tegmark discusses what consciousness could be, briefly considers Integrated information theory as a viable explanation of the phenomenon, but also enumerates the most common criticisms of the theory. "Consciousness is the way information feels when being processed in certain ways," he summarizes his own view of consciousness and speculates that it must be substrate-independent, similarly to remembering, computing and learning.

He finishes his book optimistically, describing the work he does at the Future of Life Institute he has founded. Just like his previous book, "Life 3.0" is a brilliant example of existential hope, something humanity really needs. If you want to read an enthralling, captivating book on AI, choose this one. Not that it needs any promotion after it was praised by Elon Musk :).

Follow it up by Nick Bostrom's voluminous "Superintelligence: Paths, Dangers, Strategies," another great text on the subject.

Paul says

Ask people to describe what they imagine artificial intelligence and a number of their reference points would no doubt be rooted in film and literature. There is the brutal robot from the Terminator films, the benign but deadly HAL9000 from 2001 A Space Odyssey, and the contemplative Deep Thought that Douglas Adams gave us. AI has a long way to go, but it is becoming something that people are beginning to use on a daily basis when they talk to Siri or Alexa.

The potential benefits of AI for humanity could be enormous, it could be used to run all sorts of systems, search for crimes and maybe be part of the justice process, monitor our health, assist with our jobs, and have the potential to actually do some of the most menial. People are considering using them for warfare too, one step on from what the drone does under human control at the moment.

Whilst AI excites some people who can only see the positives, after all the potential of it is huge; there are others who are very concerned that about the downsides so much so that there are AI systems that are not connected to the world wide web. Using AI for war could backfire spectacularly, bye bye human race; and what happens if the AI managing your house is hacked? Or the one in your car fails at speed. Images of those pods in the matrix come to mind...

The subjects Tegmark covers In Life 3.0 goes some way to addressing these and a lot more issues that are concerning people about the implications of AI. Some of the subjects he writes about were what you'd expect in a book like this, consciousness, intelligence, life and the implications of an AI totalitarian state, would it be a utopia or worse. There were some chapters that I didn't think were totally relevant to the subject; for example, he wanders off into the realms of space-time and goals. Was a little disappointing overall as this is a subject that needs urgent discussion right now.
