



Brain vs Machine: Human Ideas on the Age of Intelligent Machines

John Brockman

[Download now](#)

[Read Online](#) 

Brain vs Machine: Human Ideas on the Age of Intelligent Machines

John Brockman

Brain vs Machine: Human Ideas on the Age of Intelligent Machines John Brockman

As the world becomes ever more dominated by technology, John Brockman's latest addition to the acclaimed and bestselling "Edge Question Series" asks more than 175 leading scientists, philosophers, and artists: What do you think about machines that think?

The development of artificial intelligence has been a source of fascination and anxiety ever since Alan Turing formalized the concept in 1950. Today, Stephen Hawking believes that AI "could spell the end of the human race." At the very least, its development raises complicated moral issues with powerful real-world implications—for us and for our machines.

In this volume, recording artist Brian Eno proposes that we're already part of an AI: global civilization, or what TED curator Chris Anderson elsewhere calls the hive mind. And author Pamela McCorduck considers what drives us to pursue AI in the first place.

On the existential threat posed by superintelligent machines, Steven Pinker questions the likelihood of a robot uprising. Douglas Coupland traces discomfort with human-programmed AI to deeper fears about what constitutes "humanness." Martin Rees predicts the end of organic thinking, while Daniel C. Dennett explains why he believes the Singularity might be an urban legend.

Provocative, enriching, and accessible, *Brain vs. Machine* may just be a practical guide to the not-so-distant future.

Brain vs Machine: Human Ideas on the Age of Intelligent Machines Details

Date : Published October 6th 2015 by Harper Perennial (first published January 1st 2015)

ISBN :

Author : John Brockman

Format : Kindle Edition 528 pages

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

 [Download Brain vs Machine: Human Ideas on the Age of Intelligent ...pdf](#)

 [Read Online Brain vs Machine: Human Ideas on the Age of Intellige ...pdf](#)

Download and Read Free Online Brain vs Machine: Human Ideas on the Age of Intelligent Machines
John Brockman

From Reader Review Brain vs Machine: Human Ideas on the Age of Intelligent Machines for online ebook

Mark Hodges says

An expansive collection of essays concerning artificial intelligence and how it could impact our world, both from pro and con perspectives. A must read for those interested in technology and its impact on the world.

I won this book via the Goodreads giveaways

Seth Hanson says

What to Think About Machines that Think is a collection of essays by some of the most prominent scientists and experts in the field of artificial intelligence. It explores many different avenues of thought on the subject, including the morals of such a world, how we would control them, how we would treat them, what rights we would have to give them, if any and how far we would let their intelligence go. When we are capable of it, do we give them all our emotions, or is better not to burden them with the capacity to think as we do. What classifications do we use for them? Most important, is artificial intelligence even intelligence?

No one knows the answer to these questions, but the thoughts, opinions and ideas of those who can come closest have been expertly collected, recorded and compounded as a riveting account of the thought process that the greatest minds on the subject of our age take through what will probably become one of the most controversial topics of the next century.

One thing that I like about this book is the fact that its format requires it to have many differentiating opinions. Whereas other books, even those that try hard to be unbiased, are going to lean in one direction or the other, this book is able to lean in all directions at once which makes it a very interesting and encompassing read.

The only problem I have with it is that the format of short essays can get kind of old, and changes the style a lot.

I would recommend this book to anyone interested in robotics, AI or philosophy.

Matthew Geleta says

I was disappointed with this book. The idea behind the book---a collection of essays by eminent thinkers in academic fields relating to computer intelligence---is great. The final product is poor. The "essays" are stubs that are not long enough to discuss anything in depth. The stubs are often non-scientific (even pseudoscientific) and overly (at times humorously) speculative.

Jen says

This book was so enjoyable that I ended up buying it on my Kindle so I could take my time and savor it. I spent time considering each essay and evaluating if I agreed with the author and how their points fit into the larger thought structure I was forming on this topic. I loved that the authors were from such a wide pool -

artists, engineers, psychologists and everyone in between. I also really appreciate the wide range on opinions in the essays. There were essays arguing that we'd soon see a total takeover by our robot overlords and essays who firmly espoused that such a thing would never happen. In both instances, very well articulated. A highly interesting topic and a thoughtful take on it.

R K says

It was alright. Just a big book pondering the "what ifs" of AI. Fun but not really informative.

Troy Blackford says

This book holds 200 essays, and most of them are crap. The second Brockman curated 'literary symposium' that I've read that I didn't like. I've read somewhere around seven, and most of them have been brilliant. The problem is, this question largely elicited self-congratulatory moping and starry-eyed catastrophizing. Most of the respondents seemed to think we are all on the verge of being murdered by robots, and that if people don't think so, they are simpletons. Dr. Steven Pinker established an effective rebuttal to this line of thinking very early on, and thus, every subsequent argument that 'Extreme, human-like AI is near at hand' seemed ridiculous (as I agree with Pinker), and every bald argument that 'It isn't close at hand' seemed obvious. I don't need to read 200 essays about 'Why I shouldn't buy stock in a company that employs nothing but cheese and is run by a canoe,' and I didn't need to read 199 essays saying 'We either will or will not be attacked by robots.' One guy even went so far as to say 'We should remember Lord Rutherford (or something) and his pronouncement that nuclear energy would never happen: the first atomic test happened less than 24 hours later.' I thought "Okay, I'm saying you're full of shit at 8:59 PM. If I'm not killed by a robot by 9:00 PM tomorrow, you're SUPER full of shit."

But about a fifth of the essays were interesting, dealing with the topic in a more in-depth, nuanced, or uniquely-angled manner. I wish they had been specially noted, saving me the trouble of digging for them. They made the book worth it. The final note, from a trio of authors at one of Google's AI research laboratories, was the kind of rational argument the book could have used more of. At any rate, a largely tedious book with a few bright spots.

Dinesh Jayaraman says

Aside from the occasional person who really knows what (s)he's talking about, this collection of essays reads like a bunch of laypeople mouthing off about things they know very little about. Disappointing, compared to past Edge question responses.

Sean Fishlock says

What do I think about machines that think? I think that if they actually thought about it, they'd probably give this book a miss.

Mainly because it's a mess of so called experts with literally no clue and this is evident from their brief

essays in which they jump around and throw a lot of rocks, but can't agree on pretty much anything. Very little of it has any practical use from either a computer science or sociological standpoint.

I managed to group the essays into about three main categories:

1) The skeptics ... these were clearly the majority .. those stuck at the mind-body problem who still subscribe to the 1960s view that "Computers can only do what they are programmed to do". They waffle on about natural selection, millions of years of evolution to extol the wonders of the human mind in order to poo poo the idea of machines with minds even though machine intelligence has developed at a pace that wildly outstrips our own.

2) The starry eyed SETI science fiction types who speculate on bizarrely alien minds with completely own agendas, from Terminators to grey goo. Good entertainment value, but not particularly useful all the same.

3) The handful that actually seem to understand the difference between a calculator and a neural network and talked about cooperative ways advanced artificial intelligence can work with humankind.

It seemed like only the people from Google with their short and sweet epilogue sitting somewhere in between, offered something insightful for us to take away, and saved the book from its logical mess!

Eric Lawton says

A few gems among many essays that seem to have little original or useful in them.

I've read several of these Edge essay collections. This is the worst. It may be that this interesting topic is too complex to say anything useful in a page or two which is the normal length in this book. A few are much shorter, a few spill over into a third page. One of the short ones looks like the author was declining the invitation to contribute (roughly, "I don't think that machines think, so I don't have much to say").

Too many of the essays just go over the same ground. Either

"I define thinking narrowly, so based on my definition, these purported examples of machine thought don't qualify" or

"I define thinking broadly, and here are some examples of machines doing it"

Too many just give opinions of what will be possible in the future, without supporting evidence or even logic. A few at least introduce the opinion with "I suspect that", so that we know they're just gut feelings that we can ignore.

I read this on Kindle. I went back through my highlights. Only a handful were of the "interesting, think about this some more" type. Most were "here's another example of why this deserves a low rating.

Many of the thinkers, in spite of their excellent reputation in their own field, are not experts in this field and some of them are not aware of recent advances in the field. The experts generally do a good job of explaining where some marketing claims are just hype, so if you are not aware of the state of the art, you may learn something, but you'd do just as well with some Google searches.

Simón says

Not even a third through the book, I can already say I won't like it.

While the idea (a collection of essays on AI) is very good, the implementation lags behind. Most essays are a couple of pages long, don't go into any detail, and state the obvious. Some are good in that they explain the relationship between AI and statistics; or make interesting metaphors linking AI with biology and the appearance of multi-cellular organisms; or explain how natural selection will benefit AI -at some point. Others are a bunch of crap, a way to quickly put together whatever the minimum number of words was.

For a book like this to be any good, the collection should be somehow curated, filtered, even organised! There are different points of view, there are essays that go into the philosophical implications, others that talk about ethics... Why not grouping them? Why not providing an index? Maybe having a short summary before each section?

The answer is obvious: that would require extra work, which John Brockman, for whatever reason, didn't want to do. Meh.

EDIT: I have finally finished this book. I did it by skip-reading, stopping only in the essays that had something interesting to say. There weren't many, and my previous review still stands: this is way too long for how poorly curated the selection has been. In its current shape, I can only recommend to stay away from it.

Rossdavidh says

Full disclosure: back in the last millenium, I spent a little bit of time working on artificial neural networks, one of the many types of artificial intelligence. Then, as before, it was a technology of great promise, that never quite seemed to make the break into the "real world", which is to say the world of commerce rather than academia, where people are more likely to care only about its abilities, not its premise. This continued to be the case for some time thereafter, as academics and researchers in the R&D departments of sufficiently large corporations kept trying, and the results of "artificial intelligence" kept being underwhelming.

Sometime around the middle of last decade, this began to change. If one had to point to the cause, it would be tempting to say "Google", which was founded by AI researchers, and which is even more or less structured like a neural network. But more fundamentally, the cause was "Big Data". Many of the same algorithms which failed to perform in the 1990's, began to perform admirably once you gave it a million cases to learn on, instead of a hundred.

Now that AI has begun to understand human speech, identify images with faces, and otherwise do things which are useful (and perhaps even worrisomely over-useful), many of the old conversations about what "AI" really means, have acquired a new urgency. Instead of reading the musings of one particular thinker on this topic for 500+ pages, what we get in this book are the musings of over 100 thinkers, for a few pages each. Some of them you've probably heard about: Steven Pinker, Tim O'Reilly, Douglas Coupland, Brian Eno. Others, you may or may not have heard of, but should learn to listen to: Susan Blackmore, Alison Gopnik, Nicholas Christakis, Jonathan Gottschall. Nearly all of them have something interesting to say, and nearly all of it is contradicted by several others in the same collection, and their predictions range from apocalyptic to worried to cautiously optimistic to skeptical that there's even anything there to be worried or optimistic about.

Nearly every essay in this book are worth not only reading, but also putting aside and ruminating over before you move on to the next one. I mostly read it as a nightstand book, reading most often a single entry before

turning out the light, and thinking over what I had just been told, by one of the smartest, cleverest, and in most cases wisest people on the planet.

The reason that it is worth doing, of course, is that when one thinks about machines that think, inevitably one must think about how *_we_* think. Whether to explain how AI is very different, or to explain why it's destined to become very similar, it is nearly always necessary to grapple with questions such as:

- what is thinking, anyway?
- can you make a machine that thinks as well as we do, but is not conscious?
- does it matter if AI shares our values? if so, what are those, anyway?
- if I "uploaded" my thoughts, hopes, desires, memories, dreams to a machine, in a bid for Kurzweil style immortality, would that be me? How much of this tiring, sometimes painful, often (especially during allergy season) exasperating physicalness of my existence is necessary to be me? Can one be conscious without it?
- what would happen if I just became a cog in a thinking machine that was composed of all of humanity and all its machines, and it became conscious on a level I could never hope to achieve? Is that terrible? Wonderful? How do I know it hasn't already happened?

One answer to the question posed by the book's title, is that we cannot truly know what to think about our own thinking, or each other's. Are we "machines that think"? How do we know? How do I know you are (or are not)? If I can't answer that, is there any point in trying to figure out what to think about artificial thought?

That last question at least, I can answer: yes. The point of thinking about machines that think, is that it makes you think about how *_you_* think, and why. There is almost no one who could not benefit from a bit more intelligent reflection on that topic. You have the good fortune to live in a time when, for less than the cost of a single meal, you can have several hundred of the most intriguing thinkers on the planet help you ponder that. Don't waste the opportunity.

Jim Crocker says

Lots of short pieces about machine "thinking" and robots taking replacing most humans. The contributors come from astrophysics, AI, etc. Some of these guys are pretty whacky and a source of humor and insight. It makes for great reading over breakfast!

Teo 2050 says

7.5h @ 2x. This is a collection of <200, more or less thoughtful essays on "machines that think." IMO, way too many answerers hadn't really thought about all the interesting implications that could fall under this broad question. It's not just about robots walking among us (consider distributed thinking with input from all over the place), nor semantics (if "no machine will ever think," we'll have autonomous cognitive systems making decisions affecting us anyway). A minority of the answers made the book worthwhile & fun, for me, to listen to while doing mundane tasks. Many seemed to narrowly shrug the question off, unable to imagine what information technology & cognitive science will lead to even within a 100 years, not to mention long-term. *Anytime soon* is not the only time that matters. Imagine humanity('s descendants) in 2200, and it'll turn out completely different precisely because of what happens under "machines that think."

Contents:

(view spoiler)

Ed Terrell says

Where are they?" Enrich Fermi

(Intelligent machines should emerge on a relative short timescale to propagate other solar systems)

This is a great collection of short essays by great thinkers. Could machines be programmed to become sufficiently self-interested to maintain their power source? Is it perhaps a problem of vocabulary? The best essay title, for me was: "Can submarines swim" by William Poundstone. In the current state of our evolutionary language, we may not have the right words to phrase the right questions. Language must advance and adapt.

Have you ever moved 1/2way around the world by a moving company? This is how machines think. They packed bricks, they packed half drunk coke cans, they packed without thinking and they were most efficient at getting the task done. Today, spell check censors our typos. Tomorrow in China we could have "political thought" censors. Big brother, isn't the computer but the human behind the design. As one author put it: humans are cunning, and capable of deception, revenge, suspicion, and unpredictability. They are the ones to fear.

When will machines think? Or better, when will we lose our cognitive skills and we stop thinking? Our abandonment of responsibility and competence led to the global financial crisis. What is next? The complacent society? The cognitive capacity that has been freed must not be wasted. Francis Bacon spoke of "our obligation to learn, and the dream of erudition". It is up to us to decide. Obviously, we already have computer viruses that self replicate, so procreation isn't an issue. More phones are made everyday than babies are born.

The big take away is not to view intelligence anthropocentrically. We confuse agents and automata, and we will likely think that if it quacks like a duck and walks like a duck or in de Vaucanson's case it even crapped like a duck, then it must be a duck. Spontaneous arrival of self consciousness may be in the cards in some distant future but is unlikely that the Mars rover will ever become aware of itself.

Holger Matthies says

I like this series - all its books have lots of fluff, but also a handful of pearls that will stay with you and change the way you see the world. This book was no exception.

This book broke with tradition in that the majority of authors aren't expertly talking about their field, but handing out hearsay and conjecture because AI isn't their home turf.

Thus:

40% "it's going to be awesome"

30% "Skynet must stopped/won't work"

20% "But will they have rights/but what is thinking anyway"

5% "how to promote my work that has nothing to do with AI"

... and 5% are what make the book five stars. For me, the book is divided into "before and after O'Reilly", in order not to spoil. You need to read this.
