

# When Will We Reach the Singularity? – A Timeline Consensus from AI Researchers | Emerj

*Dylan Azulay*



At Emerj, we pride ourselves on presenting objective information about the applications of artificial intelligence in industry. AI is applicable in a wide variety of areas—everything from [agriculture](#) to [cybersecurity](#).

However, most of our work has been on the short-term impact of AI in business. But what about the long-term? We're not talking about next quarter, or even next year, but in the decades to come.

As AI becomes more powerful, we expect it to have a larger impact on our world, including your organization. So, we decided to do what we do best: a deep analysis of AI applications and implications. But, this time, we wanted to apply that analysis to the long-term future of AI so that we could help government and business leaders understand the more distant possibilities that this technology could bring to fruition.

We interviewed 32 PhD researchers in the AI field and asked them about the technological singularity: a hypothetical future event where computer intelligence would surpass and exceed that of human intelligence with profound consequences for society.

Throughout the course of our research, we wanted to answer six main questions:

- When will the singularity occur, if at all?
- How might post-human intelligence be most likely achieved?
- Will future artificial intelligence have a tendency to fragment, or to cluster into a singular intelligence?
- What role will humans play in a post-singularity world?

- Should international bodies like the UN play a role in guiding AI?
- What should businesses and governments do now to prepare for the future of AI?

As such, this series has a total of six chapters:

1. (You are here) **When Will We Reach the Singularity? – A Timeline Consensus from AI Researchers**
2. [How We Will Reach the Singularity – AI, Neurotechnologies, and More](#)
3. [Will Artificial Intelligence Form a Singleton – or Will There Be Many AGI Agents? – An AI Researcher Consensus](#)
4. [After the Singularity, Will Humans Matter? – AI Researcher Consensus](#)
5. [Should the United Nations Play a Role in Guiding Post-Human Artificial Intelligence?](#)
6. [The Role of Business and Government Leaders in Guiding Post-Human Intelligence – AI Researcher Consensus](#)

In the past, the singularity has been more the realm of science fiction to explore. However, many AI experts, such as [Stuart Russell](#), [Max Tegmark](#), and [Stuart Armstrong](#), take this very seriously. Our goal with this survey was to garner a grounded opinion about the singularity from dozens of AI experts, as well as their thoughts about how humanity might prepare for the future of AI.

Before we continue, we should note that this not a definitive depiction of the future; only possibilities. There was much disagreement among our experts, including whether or not the singularity would occur. But regardless, understanding the possibilities will be important for anyone in government or industry to stay ahead of the curve. As one of our respondents said:

*“The future is likely to be stranger and more unpredictable than we imagine.” – James J Hughes*

## When Will We Reach the Singularity?

If we were to ask Google Maps to navigate us to the singularity, what route would it take us? And when would we reach our destination? Our experts responded with a wide range of answers. Below is a graphical breakdown of their responses:

The first questions we asked our experts was when they expected the singularity to occur. The timeline predictions varied widely:

- 62% of respondents predict a date before 2100
- 17% of respondents predict a date after 2100
- 21% of respondents believe that the Singularity is unlikely to ever occur

Broken down a different way:

- 45% of respondents predict a date before 2060
- 34% of respondents predict a date after 2060
- 21% of respondents believe that the Singularity is unlikely to ever occur

The predictions skew towards a date sooner than 2100 or even potentially sooner than 2060. This lines up reasonably well with past polls we've done (see: [Timeline for Machine Consciousness](#)).

It's interesting to note that our number one response, "2036-2060", was followed by "likely never" as the second most popular response. Muller and Bostrom, discussed later in this article, noted in their own survey writeup that many participants who would likely fall in the "likely never" category simply didn't

respond to the survey. The researchers included a quote from a keynote speaker, Hubert Dreyfus, at the “Philosophy and Theory of AI” conference in 2011 to illustrate this point. Dreyfus said:

I wouldn't think of responding to such a biased questionnaire...I think any discussion of imminent super-intelligence is misguided. It shows no understanding of the failure of all work in AI. Even just formulating such a questionnaire is biased and is a waste of time.

In other words, we recognize and understand that like the other surveys discussed in this report, our survey may fall victim to selection bias. The people who responded to us are perhaps more likely to believe that the singularity will not only occur but will occur sooner than later. We simply can't make a judgment on whether or not this is the case.

That said, why did the experts select the dates they selected (many of which are within our lifetimes)? While our poll didn't ask for a reason for why a participant selected the timeframe they did, we suspect that a number of important AI trends may have influenced their opinions.

Many nations have jumped on the bandwagon and have released national AI strategies to promote domestic R&D and new start-ups with billions of dollars of support. Sectors from [pharma](#) to [banking](#) are becoming host to dozens of annual AI events, and hundreds or thousands of sector-specific AI vendor companies.

The clear momentum of AI funding, education, and national strategy certainly bodes well for AI's continued progress, and it's possible that these trends influenced the responses of our participants. That said, we can't be sure.

In the next three sections of the article, we compare our findings to those of two other surveys of AI researchers that were conducted in the last four years and the opinions of two prominent futurist inventors and AI thought leaders. We start with a survey from Muller and Bostrom.

## Muller and Bostrom's “Future Progress in Artificial Intelligence”

Vincent C. Muller, Professor of Philosophy at Anatolia College and President of the European Association for Cognitive Systems, and Nick Bostrom, renowned philosophy at Oxford who has written over 200 publications on superintelligence, artificial general intelligence (AGI), and other topics, [conducted a survey of AI researchers and thought leaders](#).

They asked the 550 participants four main questions, two of which we'll discuss in this article:

1. When is AGI likely to happen?
2. How soon after AGI, which they called “High-level machine intelligence,” will AGI greatly surpass human capabilities (the singularity)?

It should be noted that the questions were not phrased this way in Muller and Bostrom's survey; we have paraphrased the questions.

Muller and Bostrom reached out to four groups:

1. **PT-AI**: People who were a part of the “Philosophy and Theory of AI” conference in 2011
2. **AGI**: People who were a part of the “Artificial General Intelligence” and “Impacts and Risks of Artificial General Intelligence” conferences in 2012
3. **EETN**: Members of the Hellenic Artificial Intelligence Society
4. **TOP100**: People who were among the top 100 authors in “artificial intelligence” in terms of citations according to search results on Microsoft Academic Search

## A screen shot from page 10 of “Page 10 – Future Progress in Artificial Intelligence: A Survey of Expert Opinion”

These four groups differed slightly in the way they participate in the singularity and AGI discussion:

- PT-AI participants were mostly non-technical AI thought leaders
- AGI participants were mostly technical AI researchers
- EETN participants were comprised of only published technical AI researchers
- TOP100 participants were also mostly technical AI researchers

As previously noted, Muller and Bostrom asked participants about what they call “High-level machine intelligence,” or HLMI. They did this to mitigate bias that might have arisen had their survey included words like “artificial general intelligence” or “singularity.”

Knowing this, we have decided to substitute “HLMI” for AGI for the purposes of cohesion in this article. Muller and Bostrom define HLMI as an intelligence “that can carry out most human professions at least as well as a typical human” and refer to it as “very likely implies being able to pass a classic Turing test.” As a result, we believe our rhetorical substitution is justified.

Muller and Bostrom found some interesting results:

- Participants gave it a 10% likelihood that we would have AGI by 2022, a 50% likelihood that we'd have AGI by 2040, and a 90% likelihood that we'd have AGI by 2075. These years were all median responses.
- Participants gave it a median 10% likelihood that AGI will greatly surpass human intelligence just two years after its invention. Although they might not be confident in the two-year post-invention trajectory, they do believe that there's 75% likelihood that a superintelligence that surpasses humans will arise 30 years after the invention of AGI.
  - Perhaps worthy of note is that the AGI group of participants, mostly technical AI researchers, found it 5% more likely than the overall pool of participants that AGI would surpass humans two years (15%) after its invention and 15% more likely that it would surpass humans in 30 years (90%).
  - Strangely, the other two groups of technical researchers (EETN and TOP100) were much less likely than the AGI group to believe AGI would surpass human intelligence in even 30 years time after its invention (55% and 50% respectively).

Muller and Bostrom didn't directly ask participants when the singularity might occur, but there seems to be a significant difference between their participants and ours. Half of all of Muller and Bostrom's participants responded that there is a 90% likelihood of achieving AGI *after* 2075. The mean year was 2081.

Although there is no way to directly compare the results of our survey to the results of Muller and Bostrom's due to differing methodologies, perhaps there is something to be said about the fact that 24% of our participants, the largest reporting group, were confident the singularity would occur between 2036 and 2060, whereas 50% of Muller and Bostrom's participants gave it a 90% likelihood that we wouldn't even achieve AGI until after 2075, let alone reach the singularity.

Muller and Bostrom's participants in general thought that the singularity has a 75% chance of occurring 30 years after AGI is achieved, which puts the timeframe for the singularity several decades later than our participants believed.

## Projections from Futurists – Ray Kurzweil and Louis Rosenberg

In general, it seems our participants project a shorter timeframe for reaching the singularity than Muller and Bostrom's. This puts them more in line with [Ray Kurzweil](#), renowned futurist and Google's Director of Engineering. Kurzweil claims an 86% accuracy rate with his predictions going back to the 90s. In his [singularity timeline](#), he predicts that the singularity itself is reached by 2045.

Similarly, [Louis Rosenberg](#), PhD, inventor, and CEO and Chief Scientist of unanimous AI, has this to say about the timeframe of the singularity:

Back in the early 1990s, when I started thinking about this issue, I believed that AI would exceed human abilities around the year 2050. Currently, I believe it will happen sooner than that, possibly as early as 2030. That's very surprising to me, as these types of forecasts usually slip further into the future as the limits of technology come into focus, but this one is screaming towards us faster than ever.

Kurzweil is known for a variety of predictions about emerging technology and the [transhuman transition](#), including brain-machine interface, nanotechnologies, AGI, etc.

Kurzweil's predictions from – The Singularity is Near (2005). Source: [Wikipedia](#)

### **Grace et.al's "When Will AI Exceed Human Performance?"**

Katja Grace, John Salvatier, Allan Dafoe, Baobao Zhang, and Owain Evans, researchers from Oxford, Yale, and AI Impacts, [conducted a survey of their own](#). Their participants were 352 researchers who published at the 2015 NIPS and ICML conferences for research in machine learning. 82% of their participants were in academia, whereas 21% were working in industry.

Grace et.al use "High-level machine intelligence" (HLMI) in their survey, but unlike Muller and Bostrom, Grace et.al seem to define it the way we define the singularity: "outperforming humans in all tasks." For this reason, we are going to substitute Grace and colleagues' use of HLMI for the singularity for the purposes of cohesion in this article.

Grace and colleagues' participants gave a mean 10% likelihood that the singularity would be achieved within 9 years after the survey was conducted in 2016 (by 2025). Participants also gave a mean 50% likelihood that the singularity would occur within 45 years after the survey was conducted (by 2061). This puts their respondents on the tale end of our largest response group: those who believe the singularity will happen between 2036 and 2060.

Strangely, Grace et.al also asked participants “for the probability that AI would perform vastly better than humans in all tasks two years after HLMI is achieved.” Given their definition of “HLMI” as a machine intelligence that was “better” than humans at all tasks, it’s unclear what exactly the researchers mean by “vastly.” We’re unsure if this was a question they asked in reference to the singularity or if their definition of HLMI was in reference to the singularity like we interpreted it.

That said, the question is very similar to the one asked of Muller and Bostrom, and so we can compare the two responses. Interestingly, Grace and colleagues’ participants responded with a median 10% likelihood that AI would be “vastly” better at humans at all tasks two years after the invention of HLMI, which for this question seems to reference AGI as it does in Muller and Bostrom’s survey. This is the same likelihood that Muller and Bostrom’s participants reported.

Although it would have been interesting to see what chance Grace and colleagues’ participants gave for the singularity occurring 30 years after the invention of AGI, the researchers did not include the question on their survey.

### **So When Will the Singularity Likely Occur?**

It seems as though AI researchers and thought leaders project the singularity to occur roughly within the next four to five decades. Muller and Bostrom’s participants project a further timeline, but Grace et.al, Kurzweil, and Rosenberg all project the singularity to occur roughly within the timeframe of our largest respondent group: once again, between 2036 and 2060.

*This article was written with large contributions by Brandon Perry, who wrote the introduction and “When Will We Reach the Singularity?” subsection introduction.*