


## How Long Till Human-Level AI? What Do the Experts Say?

Ben Goertzel, Seth Baum, Ted Goertzel  
February 5, 2010  
hplusmagazine.com



When will human-level AIs finally arrive? We don't mean the narrow-AI software that already runs our trading systems, video games, battlebots and fraud detection systems. Those are great as far as they go, but when will we have really intelligent systems like C3PO, R2D2 and even beyond? When will we have Artificial General Intelligences (AGIs) we can talk to? Ones as smart as we are, or smarter?

Well, as Yogi Berra said, "it's tough to predict, especially about the future." But what do experts working on human-level AI think? To find out, we surveyed a number of leading specialists at the Artificial General Intelligence conference (AGI-09) in Washington DC in March 2009. These are the experts most involved in working toward the advanced AIs we're talking about. Of course, on matters like these, even expert judgments are highly uncertain and must be taken with multiple grains of salt — nevertheless, expert opinion is one of the best sources of guidance we have. Their predictions about AGI might not come true, but they have so much relevant expertise that we should give their predictions careful consideration.

We asked the experts when they estimated AI would reach each of four milestones:

- passing the Turing test by carrying on a conversation well enough to pass as a human
- solving problems as well as a third grade elementary school student
- performing Nobel-quality scientific work
- going beyond the human level to superhuman intelligence

We also asked how the timing of achieving these milestones would be affected by massive funding of \$100 billion/year going into AGI R&D.

We also probed opinions on what the really intelligent AIs will look like — will they have physical bodies or will they just live in the computer and communicate with voice or text? And how can we get from here to there? What kind of technical work should be prioritized? Should we work on formal neural networks, probability theory, uncertain logic, evolutionary learning, a large hand-coded knowledge-base, mathematical theory, nonlinear dynamical systems, or an integrative design combining multiple paradigms? Will quantum computing or hypercomputing be necessary? Would it be best to try to emulate the human brain as closely as possible, or are other approaches better? Should we try to make them like humans, or should we make them different? Finally, we asked the experts how we can make human-level AGIs as safe and useful as possible for the humans who will live with them.


We posed these questions to 21 AGI-09 conference participants, with a broad range of backgrounds and experience, all with significant prior thinking about AGI. Eleven of our respondents are in academia, including six Ph.D. students, four faculty members and one visiting scholar, all in AI or allied fields. Three are lead researchers at independent AI research organizations, and another three do the same at information technology organizations. Two are researchers at major corporations. One holds a high-level administrative position at a relevant non-profit organization. One is a software patent attorney. All but four participants reported being actively engaged in conducting AI research.

The detailed results of our survey will be written up for publication in the scientific literature, but we've decided to share the highlights now with h+ Magazine readers. We think you'll find them as fascinating as we did.

### **What the Experts Said About the Timing of Human-Level AI**

The majority of the experts who participated in our study were optimistic about AGI coming fairly quickly, although a few were more pessimistic about the timing. It is worth noting, however, that all the experts in our study, even the most pessimistic ones, gave at least a 10% chance of some AGI milestones being achieved within a few decades.

The range of best-guess time estimates for the AGI milestones without massive additional funding is summarized below:



Turing Test. Photo credit: wikimedia.org/Bilby


We were somewhat surprised by the ordering of the milestones in these results. There was consensus that the superhuman milestone would be achieved either last or at the same time as other milestones. However, there was significant divergence regarding the order of the other three milestones. One expert argued that the Nobel milestone would be easier than the Turing Test milestone precisely because it is more sophisticated: to pass the Turing test, an AI must “skillfully hide such mental superiorities.” Another argued that a Turing test-passing AI needs the same types of intelligence as a Nobel AI “but additionally needs to fake a lot of human idiosyncrasies (irrationality, imperfection, emotions).” Finally, one expert noted that the third grade AI might come first because passing a third grade exam might be achieved “by advances in natural language processing, without actually creating an AI as intelligent as a third-grade child.” This diversity of views on milestone order suggests a rich, multidimensional understanding of intelligence. It may be that a range of milestone orderings are possible, depending on how AI development proceeds.

One observed that “making an AGI capable of doing powerful and creative thinking is probably easier than making one that imitates the many, complex behaviors of a human mind — many of which would actually be hindrances when it comes to creating Nobel-quality science.” He observed “humans tend to have minds that bore easily, wander away from a given mental task, and that care about things such as sexual attraction, all which would probably impede scientific ability, rather than promote it.” To successfully emulate a human, a computer might have to disguise many of its abilities, masquerading as being less intelligent — in certain ways — than it actually was. There is no compelling reason to spend time and money developing this capacity in a computer.

We were also intrigued to find that most experts didn’t think a massive increase in funding would have a big payoff. Several experts even thought that massive funding would actually slow things down because “many scholars would focus on making money and administration” rather than on research. Another thought “massive funding increases corruption in a field and its oppression of dissenting views in the long term.” Many experts thought that AGI progress requires theoretical breakthroughs from just

a few dedicated, capable researchers, something that does not depend on massive funding. Many feared that funding would not be wisely targeted.

Several experts recommended that modest amounts of funding should be distributed to a variety of groups following different approaches, instead of large amounts of funding being given to a “Manhattan Project” type crash program following one approach. Several also observed that well-funded efforts guided by a single paradigm had failed in the past, including the Japanese Fifth Generation Computer Systems project. On this, one person said, “AGI requires more theoretical study than real investment.” Another said, “I believe the development of AGIs to be more of a tool and evolutionary problem than simply a funding problem. AGIs will be built upon tools that have been developed from previous tools. This evolution in tools will take time. Even with a crash project and massive funding, these tools will still need time to develop and mature.” Since these experts are precisely those who would benefit most from increased funding, their skeptical views of the impact of hypothetical massive funding are very likely sincere.



Weakness of Turing Test. Photo: [wikipedia.org/Charles Gillingham](http://wikipedia.org/Charles Gillingham)

## What Kind of Technical Approach Will First Achieve Human-Level AI?

Interestingly, we found that none of the most specific technical approaches we mentioned in our survey received strong support from more than a few experts, although the largest plurality emerged in support of probability theory. There was, however, strong agreement among the experts that integrating a wide range of approaches was better than just focusing on a single approach. There were a few who were highly bullish on robotics as the correct path to AGI in the relatively near term, whereas the rest felt robotics is probably not necessary for AGI.

## Impacts of AGI

In science fiction, intelligent computers frequently become dangerously competitors with humanity, sometimes even seeking to exterminate humanity as an inferior life form. And indeed, based on our current state of knowledge, it's hard to discount this as a real possibility, alongside much more benevolent potential outcomes. To probe this issue, we focused on the “Turing test” milestone specifically, and we asked the experts to think about three possible scenarios for the development of human-level AGI: if the first AGI that can pass the Turing test is created by an open source project, the United States military, or a private company focused on commercial profit. For each of these three scenarios, we asked them to estimate the probability of a negative-to-humanity outcome if an AGI passes the Turing test. Here the opinions diverged wildly. Four experts estimated a greater than 60% chance of a negative outcome, regardless of the development scenario. Only four experts gave the same

estimate for all three development scenarios; the rest of the experts reported different estimates of which development scenarios were more likely to bring a negative outcome. Several experts were more concerned about the risk from AGI itself, whereas others were more concerned that humans who controlled it could misuse AGI.

Several experts noted potential impacts of AGI other than the catastrophic. One predicted “in thirty years, it is likely that virtually all the intellectual work that is done by trained human beings such as doctors, lawyers, scientists, or programmers, can be done by computers for pennies an hour. It is also likely that with AGI the cost of capable robots will drop, drastically decreasing the value of physical labor. Thus, AGI is likely to eliminate almost all of today’s decently paying jobs.” This would be disruptive, but not necessarily bad. Another expert thought that, “societies could accept and promote the idea that AGI is mankind's greatest invention, providing great wealth, great health, and early access to a long and pleasant retirement for everyone.” Indeed, the experts’ comments suggested that the potential for this sort of positive outcome is a core motivator for much AGI research.

## **Conclusion**

To successfully emulate a human, a computer might have to disguise many of its abilities, masquerading as being less intelligent...

We know of two previous studies exploring expert opinion on the future of artificial general intelligence. In 2006, a seven-question poll was taken of a handpicked group of academic AI researchers (mostly not focused on AGI in their research) at the AI@50 conference. Asked “when will computers be able to simulate every aspect of human intelligence?”, 41% said “more than 50 years” and 41% said “never.” And in 2007, the futurist entrepreneur Bruce Klein gave an online survey that garnered 888 responses, asking one question: “When will AI surpass human-level intelligence?” The bulk of his respondents believed that human-level artificial intelligence would be achieved during the next half century.

In broad terms, our results concur with those of the two studies mentioned above. All three studies suggest that significant numbers of interested, informed individuals believe it is likely that AGI at the human level or beyond will occur around the middle of this century, and plausibly even sooner. Of course, this doesn’t prove anything about what the future actually holds — but it does show that, these days, the possibility of “human-level AGI just around the corner” is not a fringe belief. It’s something we all must take seriously.