

# Artificial Intelligence: Looking Forward 15 Years

Peter Stone<sup>1b</sup>, The University of Texas at Austin and Sony AI, Austin, TX, 78712, USA

*Almost 10 years ago, I co-authored a report that predicted the effects of Artificial Intelligence on daily life in the year 2030. This article reflects on and evaluates our predictions from a decade ago and looks forward another decade and a half. While there are good reasons for both excitement and apprehension, it remains within our hands, as a society, to ensure that the benefits of AI outweigh the risks.*

The whole world is abuzz due to the latest advances in artificial intelligence (AI). And rightly so. The possibilities are enormous, the dangers are real, and the dust hasn't yet settled. But when it does, AI will not have been "solved," and human beings will still be around.



Unlike in the movies, there is no race of superhuman robots on the horizon. And while the potential to abuse AI technologies must be acknowledged and addressed, their greater potential is, among other things, to make driving safer, help children learn, and extend and enhance people's lives.

## ARTIFICIAL INTELLIGENCE

Artificial Intelligence (AI) is a science and a set of computational technologies that are inspired by—but typically operate quite differently from—the ways people use their nervous systems and bodies to sense, learn, reason, and take action.

So begins the 2016 report of the One Hundred Year Study on Artificial Intelligence (AI100) that I led.<sup>a</sup> It is particularly important to understand from this definition that AI is not any single thing, but rather a collection of many different technologies. Specifically, the most recent and most powerful generative AI models, such as GPT, Gemini, Llama, and Claude, are each one of many AI-based systems, each with different capabilities, strengths, and weaknesses.

The report continues:

Though much has changed in the decade since, I still stand by these words. If you've spent any time interacting with a large language model (LLM), I suspect that you've been very impressed with its capabilities. So called *foundation models*<sup>b</sup> are able to generate text and images that are amazingly realistic. But even so, they do not come close to unlocking an understanding of the human brain, nor to fully replicating human intelligence, let alone surpassing it. There is little risk that they will soon get out of control and pose an imminent "existential" threat to humankind, at least not to the degree that nuclear weapons and pandemics already do. And the potential benefits remain enormous.

## THE NEXT 15 YEARS

The charge of the AI100 Study panel, roughly speaking, was to recap the past 15 years of progress in AI and to project the next 15 years, with regards to its impacts on daily life in large, first-world urban regions, typical North American cities. Written mostly in 2015 and 2016, it was thus entitled "Artificial Intelligence and Life in 2030" (see footnote a).

<sup>a</sup><https://ai100.stanford.edu/2016-report/>

<sup>b</sup>Generative AI models, such as ChatGPT and Gemini are sometimes called Large Language Models (LLMs), but are also more generally referred to as foundation models.

Approximately 10 years later, for this article I have been asked, once again, to look 15 years into the future. I thus reflect here on the projections in the original AI100 report, identifying some that have indeed come to pass, and others that seem to still remain 15 years (or more) in the future.

While the 2016 report was the result of extensive deliberations by 17 AI experts, in this article, I report on my own personal opinions, including some unrelated to that report.

## AI100 REVISITED

The 2016 AI100 report addressed eight different areas with likely urban impact by 2030, including two with projected relatively near-term implications (transportation and entertainment), three with medium-term implications (health care, education, and public safety and security), and three with longer term implications (service robots, low-resource communities, and the workplace). The report closes with three high-level policy recommendations. In the remainder of this section, I briefly summarize the report's conclusions in each of these areas and comment on what has happened since.

### Transportation

The report identified transportation as likely to be "one of the first domains in which the general public will be asked to trust the reliability and safety of an AI system for a critical task." While autonomous vehicles have not replaced human-driven cars, at least in some cities and on some highways, we do now share the roads with autonomous taxis and trucks.

It remains unclear to what extent, and by when, the general public will completely embrace full self-driving technology. But there has certainly been progress over the past decade. Still, transportation remains an area in which timelines continue to be pushed back such that much of what we projected for 2030 now seems reasonable as projections for 2040. The prospect of reorganizing urban centers around new forms of commuting, increased on-demand transportation, and autonomous pickup and delivery of people and packages remains futuristic and is likely to continue at a relatively slow pace.

### Entertainment

On the other hand, the main projections around entertainment in the report were that it would become more interactive, personalized, and engaging. This is indeed coming to pass already, and will surely continue to do so. The industry is still wrestling with how to balance the justifiable concerns of creators that their past content and IP be treated with respect against

the excitement for using generative AI to unlock new avenues for creativity.

GT Sophy, a superhuman racing agent in the wildly popular video game and driving simulator, Gran Turismo, initially made a splash by outperforming human champions at a real-time control task,<sup>c</sup> using an end-to-end deep reinforcement learning approach. It then went on to become one of the most widespread commercial deployments of a reinforcement learning agent by being broadly incorporated into the game.

Over the next 15 years, I expect to see video games that are built around AI technologies from the start, incorporating interactive nonplayer characters powered by modern LLMs. One way or another, we are also sure to see big impacts in the production and distribution of music, movies, and literature. And, as projected in the AI100 report, there will probably continue to be advances in targeting content to individual tastes on demand.

AI technologies continue to be transformative in the world of sports, with most major franchises embracing data analytics to build their teams and maximize the performance of their players. And we see AI-based technologies transforming the way sports are officiated and watched. This is an area in which progress has been at least as fast as projected 10 years ago, if not more so, and will likely continue to accelerate over the next 15 years.

### Health Care

Health care is often identified as the area in which AI-based technologies will provide the greatest benefits to humanity. As written in the 2016 report, "AI-based applications could improve health outcomes...for millions of people—but only if they gain the trust of doctors, nurses, and patients." Data privacy, and other regulatory and commercial obstacles remain, making practical progress in clinical settings relatively slow. But we do see increasing numbers of successful deployments, such as computer vision technologies helping in radiology and robots helping practitioners from surgeons to nurses.

Meanwhile, the implications of programs, such as AlphaFold, on drug discovery and personalized medicine are just starting to be fully understood. The promise for AI advances in health care remains enormous, and there will surely be additional major breakthroughs over the coming decade and a half.

### Education

One thing we didn't predict in 2016 was how important AI education would become. While the report focused

<sup>c</sup><https://www.gran-turismo.com/us/gran-turismo-sophy/>

on the use of AI-based technologies for education, it is now clear that education *about* AI will become essential for the next generation of workers. It may soon be required that all university graduates, no matter their discipline, be required to attain some degree of literacy in AI. With this idea in mind, my colleagues and I recently created a course called “The Essentials of AI for Life and Society” that is open to all undergraduates at the University of Texas at Austin.<sup>d</sup> Indeed, meaningful education about AI may become a fixture in elementary and secondary school education, alongside reading, writing, and arithmetic. Everyone will need to become proficient at using AI tools and at understanding what they can and cannot be trusted to do reliably.

At the same time, due to the immense and growing societal implications of AI technologies, AI students and practitioners will be required to become well-versed in ethical paradigms. As such, progress in AI may drive greater unification of science and engineering with humanities and social sciences.

While education *about* AI is a new consideration, the projections from the AI100 report on AI technologies enhancing education at all levels, while making it more broadly accessible, remain salient. Platforms, such as Khan Academy and intelligent tutoring technologies, are aggressively leveraging AI technologies to personalize curricula and democratize learning. That trend is another one that is likely to accelerate over the coming years.

## Public Safety and Security

The AI100 report highlighted opportunities and challenges for fraud detection, predictive policing, and protecting urban infrastructure. A recurring theme was the challenges that arise when it comes to avoiding the magnification of systemic, harmful biases that may be prevalent in training data, or in the way the technologies are deployed and used.

Awareness about this issue has risen dramatically over the past decade, both among practitioners and the general public. I’m optimistic that over the coming years, best practices will be established and adhered to in such a way that AI technologies will improve public safety and security.

At the same time, it cannot be denied that the risk of AI technologies being used to the *detriment* of public safety are also growing. While there is much talk about so called “guardrails,” designed to prevent AI-based technologies from being used for nefarious purposes, it is not yet clear how technically feasible it

is to make them reliable. As alluded to in the opening paragraphs of this article, I do not expect any appreciable danger of AI, on its own, posing an existential threat to humanity. But it is not clear to what extent defensive capabilities will be able to keep pace with offensive uses of AI technologies by people with evil intentions. And I do not expect the debate to subside any time soon between people who are more concerned about “AI safety” (long-term threats) and those who focus more on “AI ethics” (near-term implications) with regards to what threats are most important to safeguard against, and with what policies.

## General Purpose Service Robots

Over the past 10 years, several companies have been founded around the dream of finally delivering general-purpose service robots for homes and offices. However, the only really successful domestic robots to date remain autonomous vacuum cleaners. This is an area in which my projections for the next 15 years are quite similar to those from 10 years ago.

Powerful new sensors and actuators continue to be developed, and we are even starting to see impressive humanoid robots at decreasing prices. However creating reliable, affordable hardware remains very difficult, with the challenges of bringing products to market often being underestimated, leading to limited commercial opportunities in this space in the near future.

For this situation to change, there will need to be significant breakthroughs toward providing reliable, low-cost robot manipulation, especially *mobile* manipulation. And significant challenges remain related to ethics and privacy issues for robots in homes and public spaces.

## Low-Resource Communities

Not surprisingly, I know of relatively little progress that has been made over the past decade when it comes to leveraging AI technologies for helping low-resource communities, such as people without homes, rural populations, and third-world countries. Of course the greatest factor is funding. While the race for AI dominance remains as heated as it currently is among the largest international technology companies, it is unlikely that significant resources will be devoted to more philanthropic outlets in the near future. The greatest hope is probably that general-purpose technology can be easily repurposed with little to no special-purpose development required.

## Workplace

The 2016 AI100 report concluded that in the near term, AI technologies would replace more tasks than jobs,

<sup>d</sup><https://www.cs.utexas.edu/pstone/Courses/109fall23/>

while also creating new types of jobs that could be difficult to foresee—that the fear of replacing all human jobs in one generation were overblown. But it warned that the gap between the rich and the poor could grow.

Broadly speaking, these projections appear to have been accurate. Few, if any, jobs have disappeared entirely. But many jobs have been transformed, with AI-literate employees having distinct advantages over employees who have been slow to familiarize themselves with how and when to use AI-based tools. And new jobs, such as *Prompt Engineer*, definitely were not foreseen in 2016. These effects are likely to become even more emphatic as general AI capabilities continue to be transformed into industry-specific software tools.

The concern about AI-related wealth being concentrated in too few hands remains.

## Public Policy

Finally, the 2016 AI100 report made three concrete recommendations pertaining to public policy:

- Define a path toward accruing technical expertise in AI at all levels of government.
- Remove the perceived and actual impediments to research on the fairness, security, privacy, and social impacts of AI systems.
- Increase public and private funding for interdisciplinary studies of societal impacts of AI.

These recommendations remain salient, although positive progress has already been made on all three.

As indicated in the “[Education](#)” section, more and more people are being educated about AI, leading to opportunities for AI-literate personnel in all sectors of the economy, including government. Educational programs at scale, such as the University of Texas at Austin’s online Masters in AI, are contributing significantly.

Meanwhile, the level of nuance across disciplines, and even in the press and among the general public, has been elevated considerably. I expect that this positive development will be maintained over the coming years.

## UNDERSTANDING GENERATIVE AI

One unanticipated research area that has emerged over the past 10 years is the science pertaining to understanding how generative AI models represent information. Until now, a lot of the focus has been on understanding what LLMs and other generative AI models (e.g. VLMs) can and cannot do. What capabilities are currently exhibited? What capabilities will be exhibited by simply scaling up the amount of data

and compute (albeit at enormous cost)? I think of this sort of study as being analogous to neuroscience and psychology studies of the human brain, and expect interesting progress to continue over the coming years.

While there are still new architectural innovations being introduced and studied, we are now seeing a shift toward investigating ways to embed generative AI models in broader systems. Some people are calling these “agentic” systems, which includes LLMs calling external tools and/or interacting with each other in multi-LLM systems. There are emerging interesting approaches that are starting to resemble “cognitive architectures,” as they were referred to in the past, with LLMs playing a relatively limited role within a larger system that includes perception, memory, planning, and other functionality, rather than trying to use them to do everything. I am very interested to see how such embeddings of generative AI in broader systems evolves over the coming years.

## SUPERHUMAN AI?

Finally, it is worth addressing the extent to which AI technologies will achieve superhuman capabilities in the coming decade and a half.

Over the past 75 years, advances in AI have repeatedly threatened the “specialness” of humankind. Just as advances in astronomical knowledge toppled the notion that our planet is at the center of the universe, and the theory of evolution poked holes in the notion that humans are fundamentally distinct from other animals, AI landmarks have threatened the notion that we are the most intelligent entities in the universe.

Considering only events during my own research career over the past 30 years, I recall, or have been directly involved in, the demonstrations of superhuman AI systems in Chess (1998), Checkers (2007), Jeopardy! (2011), Go (2015), Poker (2017), DOTA 2 (2019), high-speed simulated racing in Gran Turismo (2022), and real-world drone racing (2023), among others. In 2024, surprisingly strong performance was exhibited on problems from the International Mathematics Olympiad. Other tasks still remain beyond the reach of AI systems, including a few games (e.g. bridge) and most sports (e.g. table tennis and soccer).

I have recently served as the president of the RoboCup Federation, an international organization that was originally motivated by the goal of creating a team of superhuman robot soccer players (and has since expanded to embrace other related goals as well). While I expect people to remain better than robots at soccer over the next 15 years (the target date for RoboCup has always been 2050), other tasks are sure to “fall”



over the next 15 years. Which will they be? Will they be mostly games and online, purely software challenges? Or will we start to see more demonstrations of superhuman capabilities in physical tasks, such as safe driving, real-world sports, picking and placing, folding laundry, cooking, and other tasks leading to the (so far) elusive dream of *general purpose service robots*?

I don't purport to know for sure which tasks will prove to be most amenable to full automation. But I do suspect that games and online tasks will remain considerably more prone to AI systems outperforming the very best people compared to real-world tasks.

### WHAT CAN/SHOULD WE DO?

In my opinion, both the concerns and the optimism surrounding AI are well-justified. AI technologies, like many other technologies, can be used for nefarious purposes. But the potential positive impacts in areas, such as health care, are enormous.

I'm encouraged that AI is becoming a very interdisciplinary area, with humanists and social scientists, including economists and political scientists, studying its societal implications deeply, in close collaboration with computer scientists and engineers. Governments are crafting policies and regulations with vigorous public debate, and with a greater degree of information about the technologies than they've had in the past.

I am personally involved in organizations, such as the AI100 and Good Systems at the University of Texas at Austin<sup>e</sup> that actively convene people with these differing perspectives. And many companies are maintaining AI ethics teams who work closely with their AI scientists and engineers, including at Sony AI, where we have a world-class AI Ethics group. I hope that these sorts of efforts and organizations continue to gain prominence.


Of course, there remains much uncertainty regarding how AI will progress over the coming 15 years. But through education and ongoing interdisciplinary research and collaboration, I'm optimistic that collectively, we will be able to shape the future of AI progress and deployment toward exhibiting a net positive influence on society.

<sup>e</sup><https://bridgingbarriers.utexas.edu/good-systems>

Over the next 15 years, AI technologies will continue to both improve educational outcomes and, increasingly, AI itself will become the focus of broad interdisciplinary education. It will also continue driving innovations in health care, transportation, entertainment, and several other related areas, as discussed throughout this article. But while seeking to realize the immense positive possibilities of AI-based technologies, we must also remain cognizant of the potential downsides. It is up to us, as a society, to shape the ways in which AI will affect how the next 15 years, and beyond, will unfold.

**PETER STONE** is a professor and director of the Learning Agents Research Group in the Department of Computer Science at the University of Texas at Austin. He is also the director of Texas Robotics, chief scientist of Sony AI, and past-chair of the AI100 Standing Committee. His research focuses mainly on machine learning, multiagent systems, and robotics. Contact him at [pstone@cs.utexas.edu](mailto:pstone@cs.utexas.edu).

## stay on the Cutting Edge of Artificial Intelligence



**IEEE Intelligent Systems** provides peer-reviewed, cutting-edge articles on the theory and applications of systems that perceive, reason, learn, and act intelligently.

