

5 years: Widespread deployment of next-generation middleware for AI and robotics and accompanying testing standards for real-world conditions.

10 years: Industry adoption of robust security and safety standards for AI applications built on common middleware.

20 years: Rich ecosystem of user-centered and privacy-preserving AI applications that are portable across different robot platforms.

4. Major Findings

Several major findings from community discussions drive the recommendations in this Roadmap:

I. Enabled by strong algorithmic foundations and propelled by the data and computational resources that have become available over the past decade, AI is poised to have profound positive impacts on society and the economy. AI has become a mature science, leveraging large datasets and powerful computing resources to produce substantial progress in many areas: exploration and training of statistical models, for instance, and powerful image and video-processing techniques. Many other areas of AI might be amenable to the same dramatic leaps forward, but are starving for appropriate data. And while collection, processing, and annotation of data are key aspects of an experimental science, architectures and frameworks are also instrumental in AI solutions. Enabled by the right theoretical and applied foundations and fueled by massive datasets and growing computational power, future AI-driven successes could affect many aspects of society, including healthcare, education, business, science, government, and security (as shown in Figure 3): removing humans from harm's way in dangerous yet vital occupations; aiding in the response to public health crises and natural disasters; expanding educational opportunities for increasingly larger segments of our society; helping local, state, and federal government agencies deliver broader range of valuable services to their citizens; or providing personalized lifelong health- and wellness care accessible to each individual's changing needs.

II. To realize the potential benefits of AI advances will require audacious AI research, along with new strategies, research models, and types of organizations for catalyzing and supporting it. Over its first few decades, AI research was characterized by steady progress in understanding and recreating intelligent behaviors, with deployments of AI-enhanced systems in narrow application areas. Since the mid-1980s, fundamental AI research has been supported largely by short-term grant-funded projects in small single-investigator labs, limiting the types of empirical work and advances possible. In recent years, the experimental possibilities enabled by data and computationally resource-rich and generously staffed industry labs have yielded significant advances, enabling wide deployment of AI-enabled systems in many societally relevant and important venues. Cross-fertilization with other fields, ranging from social science to computer architecture, is also critical to modern AI, given the demands, breadth, and implications of its applications. To adequately address the next generation of AI challenges will require sustained effort by large, interdisciplinary teams supported by appropriate resources: massive datasets, common architectures and frameworks, shared hardware and software infrastructure, support staff, and sustained, long-term funding. This Roadmap offers new models for the resources, the critical mass, and the long-term stability that will be needed in order to enable a new era of audacious AI research that is significantly more integrative and experimental while also recognizing the need for caution regarding the impact of AI in society.

III. The needs and roles of academia and industry, and their interactions, have critically important implications for the future of AI. Building on the foundations of past AI research, most of which was conducted in academia, the private sector has compiled and leveraged massive resources—datasets, knowledge graphs, special-purpose computers, and large cadres of AI engineers—to propel powerful innovations. These assets, which provide major competitive advantages, are generally proprietary. Furthermore, the constraints, incentives, and timelines in these two communities are very different: Industry is largely driven by practical, near-term solutions, while academia is where many of the fundamental long-term questions are studied. Solutions to

the next generation of AI problems cannot come from either academia or industry alone. Without the right resources, academic AI research is limited; without answers to foundational questions, industrial AI innovations will be limited. The ability to explore different approaches and models through experiments on practical problems in real-world settings is limiting the ability of universities to do experimental research.

IV. Talent and workforce issues are undergoing a sea change in AI, raising significant challenges for developing the talent pool and for ensuring adequate diversity in it. There is great demand in industry for AI talent and the gap between the supply and the demand is likely to grow significantly over the next decades. All US universities are looking for AI faculty—and struggling to hire, particularly at senior levels and in areas relevant to industry needs. Students flock to AI course offerings and research groups. Faculty across a broad swath of fields (including humanities as well as social and physical sciences, and in such professional studies as law, medicine, and public health) seek AI collaborators, and new data-science programs continue to emerge. Drawn by the higher salaries, extensive data and infrastructure resources, and numbers of potential collaborators and support staff available in industry, many AI faculty are leaving academia or pursuing appointment arrangements where they split time between a university and a corporate research lab. While this trend creates opportunities for companies and universities to engage in innovative ways of working together to advance AI, it also has the potential to negatively affect AI education, training, and the AI research pipeline. Diversity is also a major issue. Although the number of women in STEM fields is increasing, the number of women in computer science has halved since the 1980s. In AI, the participation by women and underrepresented minorities is even lower than in computing as a whole, despite the increases in college graduation rates in these groups.

V. The rapid deployment of AI-enabled systems is raising serious questions and societal challenges encompassing a broad range of capabilities and issues. Novel AI technologies have been deployed quickly into market before incorporating systems engineering, safety design principles, best practices, and societal considerations. Research is needed into the incorporation into AI systems of responsibility criteria in autonomy and assistance in order to enable frameworks that are suitable for operations in critical domains, as well as a clear articulation of the limits of AI-enabled systems—settings where such systems might on balance provide more harm than good. Research is also needed into what these best practices might be and what kinds of institutional or legal scaffolding would make them most effective. It is critical for these efforts to be supported as an integral part of AI research and development, rather than as ex post facto analysis of systems after they are built. Considerations include the fairness of decisions made by these systems, as well as their potential to introduce and amplify structural patterns of inequality and distortions of the truth. The ethical ramifications of AI-assisted decision making and content generation are critical near-term issues, given that these technologies are being used as replacements of, or assistants to, human decision makers in areas as crucial to individual lives and societal well-being as criminal justice, predictive policing, credit risk systems, employment, college admissions, student and teacher evaluation, autonomous vehicles, and national security. In all of these diverse settings, AI technologies have the potential to introduce profound structural change in the way we make decisions, who gets to make them, and for whom. AI assistance is also being used extensively for content filtering and dissemination—especially on social media, mostly invisibly, and sometimes with significant consequences. It is important to ensure that these systems are designed and deployed safely and responsibly, and with adequate oversight. It is presently challenging to explain the decisions of current AI methods, for instance, leaving the user in the dark as to which factors influenced the outcomes. Addressing this constellation of challenges will require deep, extensive studies of the interactions among AI, ethics, and society, with concern for a broad range of normative questions including fairness, accountability, and transparency in AI-assisted decision making. We will also need to identify different ways of thinking about the design of these systems so that such concerns are taken into account before they are deployed and cause harm. This kind of multidisciplinary effort is naturally and best located at the nexus of the social sciences, humanities, and computing and within university environments. Locating these efforts within university environments is also important in the light of rapidly differing incentives between industry and academia that are influencing the increasing tension around the social and ethical ramifications of widely deployed AI systems.

Substantial Long-Term Investments Drive Significant Advances

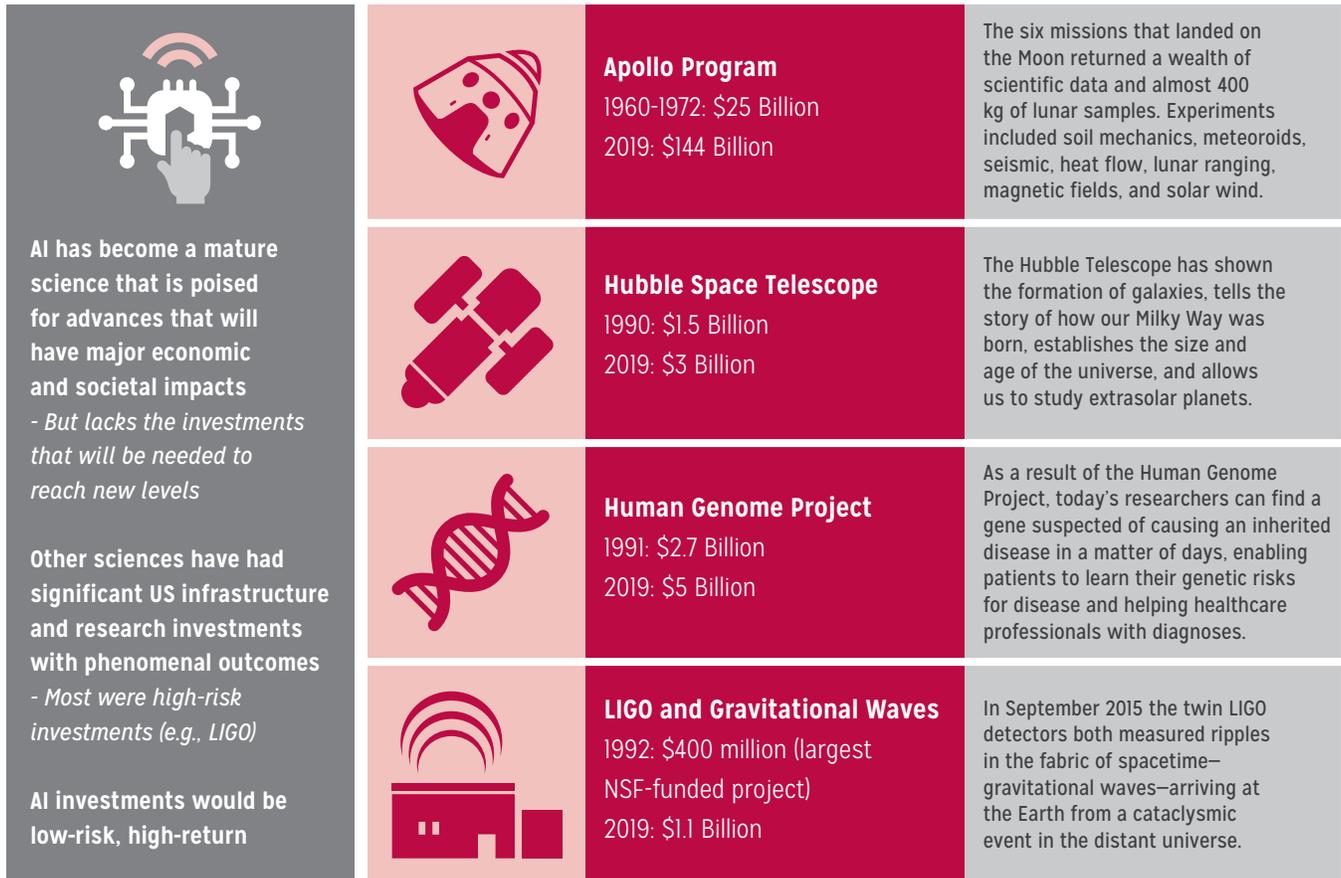


Figure 4. Substantial Long-Term Investments Drive Significant Advances

VI. Significant strategic investments in AI by the United States will catalyze major scientific, technological, societal, and economic progress. Important advances have been propelled by significant US investments in audacious projects over the past 50 years. The Laser Interferometer Gravitational Wave Observatory (\$1.1B), for instance, led to the discovery of gravitational waves. The Human Genome Project (\$2.7B) was the driver of major innovations in biomedical research. The Apollo program (\$144B) not only accelerated space travel but also catalyzed many scientific contributions. AI is ready for similar forward leaps. Recognizing this, other major industrialized countries are already embarking on major AI research and education programs. As of the time

Figure 4 References:

Apollo Program: <https://airandspace.si.edu/explore-and-learn/topics/apollo/apollo-program/>

Hubble Space Telescope: <https://www.nap.edu/read/11169/chapter/5#20>

Human Genome Project: <https://report.nih.gov/nihfactsheets/viewfactsheet.aspx?csid=45>

LIGO and Gravitational Waves: <https://www.ligo.caltech.edu/detection>

of this writing, Germany²³ and France²⁴ have allocated 3B and 1.5B euros to AI, respectively. The UK²⁵ has pledged an investment of 1B pounds in AI, together with dedicated funding for 1,000 PhDs and 8,000 specialized teachers in AI, and has repurposed its flagship Turing Institute into a major data-driven AI research center. China²⁶ has announced that it will invest billions in AI over the next five years, creating at least four \$50M/year AI Centers and a \$1B/year National AI Research laboratory with thousands of AI researchers and engineers, and committing to training 500 instructors and 5,000 students at major universities. Significant investments in AI, structured with the guidance of this document, will allow the US to take the forefront in propelling the field into a new research era and create significant impact across all sectors of society and the economy.

5. Recommendations

The progress desired in integrated intelligence, meaningful interactions, and self-aware learning pose substantial research challenges that require significant investments that are sustained over time. History tells us that there is no silver bullet, no substitute for sustained effort. For example, speech recognition took four decades of research, mostly in academia, until the work was mature enough to be commercialized (e.g., the commercial Dragon speech recognition software first appeared in the market in the 1990s) and eventually mainstream (e.g., conversational assistants). Similarly, the basic ideas in neural networks were first developed in the 1980s, and it took decades of research as well as improvements in computational resources before they were able to take off. This research Roadmap is designed to prioritize investments and foster similar advances in technologies across AI over the next two decades.



Figure 5. Overview of the 20 Year AI Research Roadmap Recommendations: A national AI infrastructure combined with training of an all-encompassing AI workforce.

²³ "Germany plans 3 billion in AI investment: government paper," Reuters, November 13, 2018, <https://www.reuters.com/article/us-germany-intelligence/germany-plans-3-billion-in-ai-investment-government-paper-idUSKCNII1AP>.

²⁴ "Emmanuel Macron Talks to Wired about France's AI Strategy," Wired, March 31, 2018, <https://www.wired.com/story/emmanuel-macron-talks-to-wired-about-frances-ai-strategy/>.

²⁵ "U.K. Unveils \$1.4 Billion Drive Into Artificial Intelligence," Bloomberg, April 25, 2018, <https://www.bloomberg.com/news/articles/2018-04-25/u-k-announces-1-4-billion-drive-into-artificial-intelligence>.

²⁶ "Understanding China's AI Strategy," Center for a New American Security, February 6, 2019, <https://www.cnas.org/publications/reports/understanding-chinas-ai-strategy>.