1. Introduction

This Roadmap articulates the potential of artificial intelligence (AI) research advances to significantly improve our world and benefit our lives. It is the result of a community activity to formulate AI research priorities for the next 20 years, and make recommendations in support of their pursuit. This research agenda is motivated by a broad range of societal drivers aiming to improve health, education, science, innovation, justice, and security. The recommendations are supported by major findings resulting from the community discussions.

The rest of this chapter provides a brief background on AI and an overview of the community activities that led to this Roadmap document. Chapter 2 (page 12) gives an overview of major societal drivers, emphasizing the potential value of AI innovation in different sectors. Chapter 3 (page 16) details the outcomes of the three community workshops held on the core technical areas of the AI research Roadmap, including specific vignettes that describe AI systems for those societal drivers and articulates the research required that spans a wide range of areas of AI. The major findings and recommendations that resulted from reflecting on those efforts are detailed in Chapter 4 (page 81) and Chapter 5 (page 84).

1.1 What Is Artificial Intelligence?

In this report, we adopt the definition from the One Hundred Year Study of AI[1] that views AI as "a branch of computer science that studies the properties of intelligence by synthesizing intelligence." The bolding is ours:

> Notably, the characterization of intelligence as a spectrum grants no special status to the human brain. But to date ***human intelligence*** has no match in the biological and artificial worlds for sheer versatility, ***with the abilities "to reason, achieve goals, understand and generate language, perceive and respond to sensory inputs, prove mathematical theorems, play challenging games, synthesize and summarize information, create art and music, and even write histories."***[2]
>
> This makes human intelligence a natural choice for benchmarking the progress of AI. It may even be proposed, as a rule of thumb, that any activity computers are able to perform and people once performed should be counted as an instance of intelligence. ***But matching any human ability is only a sufficient condition, not a necessary one. There are already many systems that exceed human intelligence, at least in speed,*** such as scheduling the daily arrivals and departures of thousands of flights in an airport.
>
> [...]
>
> AI can also be defined by what AI researchers do. ***This report views AI primarily as a branch of computer science that studies the properties of intelligence by synthesizing intelligence.***[3] Though the advent of AI has depended on the rapid progress of hardware computing resources, the focus here on software reflects a trend in the AI community. More recently, though, progress in building hardware tailored for neural-network-based computing[4] has created a tighter coupling between hardware and software in advancing AI.

---

1 "One Hundred Year Study on Artificial Intelligence (AI100)," Stanford University, accessed August 1, 2016, https://ai100.stanford.edu.

2 Nils J. Nilsson, The Quest for Artificial Intelligence. New York: Cambridge University Press, 2009.

3 Herbert A. Simon, "Artificial Intelligence: An Empirical Science," Artificial Intelligence 77, no. 2 (1995):95–127.

4 Paul Merolla, John V. Arthur, Rodrigo Alvarez-Icaza, Andrew S. Cassidy, Jun Sawada, Filipp Akopyan, Bryan L. Jackson, Nabil Imam, Chen Guo, Yutaka Nakamura, Bernard Brezzo, Ivan Vo, Steven K. Esser, Rathinakumar Appuswamy, Brian Taba, Arnon Amir, Myron D. Flickner, William P. Risk, Rajit Manohar, and Dharmendra S. Modha, "A Million Spiking-Neuron Integrated Circuit with a Scalable Communication Network and Interface," Science, vol. 345, no. 6197 (August 8, 2014): 668-73. https://science.sciencemag.org/content/345/6197

"Intelligence" remains a complex phenomenon whose varied aspects have attracted the attention of several different fields of study, including psychology, economics, neuroscience, biology, engineering, statistics, and linguistics. Naturally, the field of AI has benefited from the progress made by all of these allied fields. For example, the artificial neural network, which has been at the heart of several AI-based solutions[5,6] was originally inspired by thoughts about the flow of information in biological neurons.[7]

AI research is sometimes focused on mimicking human intelligence. For example, one approach to learn from data is to design artificial neural networks that are composed of neuron-like units that adjust how they fire based on stimuli (the data). In many cases the approaches developed are engineered for performance and do not have any resemblance to human intelligence. For example, chess playing programs consider many more possible moves than humans are able to do.

Although AI has been a field of study for several decades, the term "AI" has become a colloquial term that is used very loosely.

Sometimes "AI" is equated with machine learning, and specifically with learning from large amounts of data in order to make predictions. Although learning from data is an important area of research in AI, there are many other learning techniques that are studied in AI, including learning from tutorial instruction, learning from demonstrations, learning from a single example, learning to work in teams, etc. And although learning is ubiquitous in intelligent behaviors, machine learning is one of many areas of study in AI, such as understanding and generating language, acting based on sensor inputs, representing knowledge in machine-readable structures, coordinating with others to work in teams, and reasoning about complex problems to generate solutions.

Sometimes "AI" is used to describe robots and other computer systems that have a physical dimension such as intelligent sensors. These AI systems can operate independently or be governed by humans. Many AI systems lack a physical dimension, such as conversational interfaces and intelligent search or recommendation systems.

Sometimes "AI" is used to describe very simple programs that give the illusion of intelligence, such as "game AIs" that govern synthetic players in computer games. In these cases, the meaning is simply "a computer program that generates behaviors that are more interesting than traditional ones." A consequence of this is that AI is a moving target: what is considered AI today may not be considered AI tomorrow as more sophisticated programs are created.

"AI" is often used to refer to systems that exhibit any intelligent capability or aspect of intelligence. For example, any computer system that is able to learn from some data, or solve a problem, or generate language, is often called "the AI."

This document uses the term "AI system" to refer to a computer system that exhibits some intelligent behavior, and the surrounding text will be clear on what the intelligent behaviors are. It is important to note that just because a system includes some aspect of intelligence, it does not mean that it has general intelligent capabilities. In fact, there are no AI systems today that have general intelligent capabilities. This document lays out a Roadmap for AI research in the next 20 years, and even at that point AI systems will be far from having general intelligent capabilities.

---

[5] Gerald Tesauro, "Practical Issues in Temporal Difference Learning," Machine Learning, vol. 8, issue 3-4 (May 1992): 257-77.

[6] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis, "Mastering the game of Go with deep neural networks and tree search," Nature 529 (2016): 484-89.

[7] Warren S. McCulloch and Walter Pitts, "A logical calculus of the ideas immanent in nervous activity," Bulletin of Mathematical Biophysics, 5 (1943): 115-33.

1.2 AI Roadmap Process

As part of its continuing efforts to advance computing research, the Computing Community Consortium (CCC) reached out to the leadership of the Association for the Advancement of Artificial Intelligence (AAAI) to spearhead a 20-year Roadmap for AI.

The mission of the Computing Research Association's CCC is to catalyze the computing research community and enable the pursuit of innovative, high-impact research. CCC conducts activities that strengthen the research community, articulate compelling research visions, and align those visions with pressing national and global challenges. CCC communicates the importance of those visions to policymakers, government and industry stakeholders, the public, and the research community itself.

AAAI is a nonprofit scientific society devoted to advancing the scientific understanding of the mechanisms underlying thought and intelligent behavior and their embodiment in machines. AAAI also aims to increase public understanding of artificial intelligence, improve the teaching and training of AI practitioners, and provide guidance for research planners and funders concerning the importance and potential of current AI developments and future directions. AAAI has more than 200 elected fellows, and over 4,000 members.

Yolanda Gil (AAAI President; University of Southern California) and Bart Selman (AAAI President Elect; Cornell University) led the effort and designed a series of three community workshops:

- **Integrated Intelligence**, November 14-15, 2018, co-led by Marie desJardins (Simmons University) and Ken Forbus (Northwestern University).
- **Meaningful Interactions**, January 8-9, 2019, co-led by Kathy McKeown (Columbia University) and Dan Weld (University of Washington)
- **Self-Aware Learning**, January 17-18, 2019, co-led by Tom Dietterich (Oregon State University) and Fei-Fei Li (Stanford University)

Altogether, 91 researchers from 41 academic institutions and nine industrial partner organizations took part in this effort. Detailed participant lists appear in the Appendix.

Workshop participants discussed the potential benefits of AI for several societal drivers: reducing healthcare cost, accelerating scientific discovery, facilitating universal personalized education, innovation for business, and evidence-driven social opportunity. The workshop agendas combined full-group discussions and breakout sessions, some centered on the societal drivers and some centered on the technical challenges in the topical area. Participants discussed research priority areas and formulated 5-year, 10-year, and 15-year milestones toward 20-year stretch goals. Workshop participants also reflected on related topics such as the historical evolution of the field of AI, the multidisciplinarity of the AI research agenda, AI research infrastructure, the rapid growth of AI in industry labs, the interplay between academic and industry research, AI funding programs and agencies, and the public understanding and perception of AI technologies. Following the workshops, individual reports were drafted based on participant discussions at each event, then edited by the co-leaders and circulated to the workshop participants for comments. Findings and recommendations brought up in discussions throughout the workshops were also drafted.

A preliminary version of the Roadmap was presented to the community in a town hall meeting at the AAAI 2019 Conference on January 28, 2019, which was attended by more than 500 participants. A draft of the Executive Summary was posted on the CCC website on March 13, 2019, for preliminary comment. Numerous suggestions were collected during those months and incorporated into the Roadmap. The findings and recommendations were further refined.

A full draft of the report, including the findings, recommendations, individual workshop reports, and supporting materials, was distributed on May 13, 2019, for community comment.

1.3 Why Now?

Decades of research in AI have produced formidable technologies that are providing immense benefit to industry, government, and society. The bulk of AI research in industry at present is driven by advertising and product recommendation systems. The flood of recent developments in this area—catalyzed by the availability of massive datasets, vast knowledge graphs, powerful special-purpose computers, and large cadres of AI engineers—have created a trillion-dollar industry that is projected to quadruple in three years. While AI solutions have the potential for transformative impacts across all sectors of society and the economy, there are concerns about the security and vulnerability of these systems. As the resources that are driving the AI revolution continue to grow, the development and deployment of these technologies is poised not only to continue, but to accelerate. For all these reasons, the time is now to undertake a thoughtful and comprehensive envisioning of interdisciplinary, forward-looking R&D that will drive scientific and economic advances in AI while taking into consideration issues around security, vulnerability, policy, and ethics.

1.4 How to Read This Report

While the entire report is important for understanding and carrying out the goals of the Roadmap, we realize that it is long. All of the sections in this report are written to be read independently, so readers can jump around and read individual sections of their choosing. On your first read, it is suggested that you begin with the executive summary followed by the introduction, major findings, recommendations, and then conclusions. On your next read, please go back and read the major societal drivers and core technical areas.

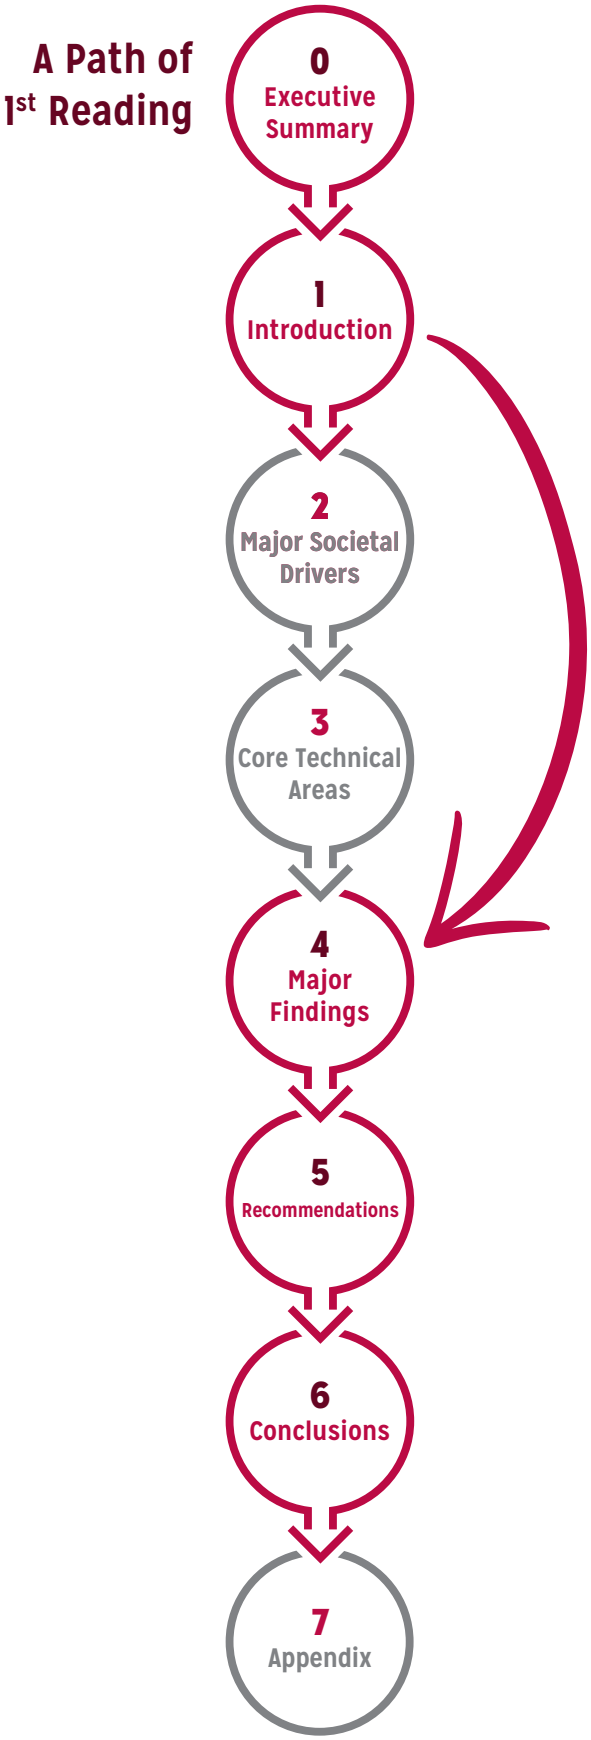

Figure 1. How to read the report

1.5 Reference Documents

We refer to the following documents for additional information about AI and its current context:

- **The state of the art in AI:** The "One Hundred Year Study on Artificial Intelligence"[8] provides a thorough review of the state of the art in AI research and applications, with contributions from researchers across all areas of AI and related fields. The report reviews the state of the art in a diversity of areas of AI research and their relevance to a range of applications with societal benefits of AI. The report provides citations and references to the published literature. This Roadmap document builds on and extends this study by providing a more in-depth motivation and description of areas for future research.
- **US government programs and priorities in AI:** The "National Artificial Intelligence R&D Strategic Plan"[9] was developed by the Networking and Information Technology Research and Development (NITRD) Program through its Task Force on Artificial Intelligence. The report reviews existing programs that fund AI research, identifies strategic priorities, and makes recommendations to guide federally funded AI research. This Roadmap document is consistent with the strategies proposed in the NITRD plan, and extends it with more specific recommendations and more elaborate descriptions of research priorities.
- **Related research Roadmaps:** "A Roadmap for US Robotics"[10] gives a thorough treatment of research priorities and applications in diverse areas of robotics and cyberphysical systems. This document is consistent with that report, and gives more extensive details on the cognitive capabilities needed in intelligent robots as well as expanding on many aspects of AI systems that inhabit the digital rather than the physical realm.

2. Major Societal Drivers for Future Artificial Intelligence Research

AI has permeated our lives and has become an engine for innovation across all sectors of society. Government investments can have a profound impact and transform society for the betterment of all citizens. Below we describe six societal areas that will be transformed by AI. These are meant to be examples of AI's potential impact to society, but in no way are they an exhaustive list. Detailed discussions for each of the societal drivers are included in the workshop reports through a series of vignettes that describe how a person or organization could interact with AI systems or would benefit from AI technologies, and discuss the capabilities that would be required and that motivate the AI research in this Roadmap. By looking at and working on the research challenges grounded in AI for social good, we will not only be changing society, but we can motivate K-12 students to study AI by presenting exciting problems to be worked on. These same societally relevant problems will help to engage women and other underrepresented populations in AI research, creating a more diverse workforce to better tackle the problems. In highlighting AI for social good, we will dispel misconceptions about AI while solving real-world, tangible problems.

2.1 Enhancing Healthcare and Quality of Life

Although the potential benefits for AI in healthcare have been demonstrated, this technology is largely untapped in clinical settings. AI applications can now diagnose some skin cancers more accurately than a board-certified dermatologist, and do so faster and

---

[8] "One Hundred Year Study on Artificial Intelligence (AI100)," Stanford University, accessed August 1, 2016, https://ai100.stanford.edu.

[9] "The National Artificial Intelligence Research and Development Strategic Plan," National Science and Technology Council, October 2016, https://www.nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf.

[10] "A Roadmap for US Robotics: From Internet to Robotics, 2016 Edition," October 31, 2016, https://cra.org/ccc/wp-content/uploads/sites/2/2016/11/Roadmap3-final-rs-1.pdf