

A 20-Year Community Roadmap for Artificial Intelligence Research in the US

Roadmap Co-chairs:

Yolanda Gil, AAAI President / University of Southern California

Bart Selman, AAAI President Elect / Cornell University

Workshop Chairs:

Tom Dietterich, Oregon State University

Ken Forbus, Northwestern University

Marie desJardins, Simmons University

Fei Fei Li, Stanford University

Kathy McKeown, Columbia University

Dan Weld, University of Washington

Steering Committee:

Elizabeth Bradley, CCC Vice Chair / University of Colorado Boulder

Ann Schwartz Drobni, CCC Director

Mark D. Hill, CCC Chair / University of Wisconsin-Madison

Daniel Lopresti, Lehigh University

Maja Mataric, University of Southern California

David Parkes, Harvard University

August 2019

Sponsored by the Computing Community Consortium (CCC)
and the Association for the Advancement of Artificial Intelligence (AAAI)



Executive Summary

Decades of research in artificial intelligence (AI) have produced formidable technologies that are providing immense benefit to industry, government, and society. AI systems can now translate across multiple languages, identify objects in images and video, streamline manufacturing processes, and control cars. The deployment of AI systems has not only created a trillion-dollar industry that is projected to quadruple in three years, but has also exposed the need to make AI systems fair, explainable, trustworthy, and secure. Future AI systems will rightfully be expected to reason effectively about the world in which they (and people) operate, handling complex tasks and responsibilities effectively and ethically, engaging in meaningful communication, and improving their awareness through experience.

Achieving the full potential of AI technologies poses research challenges that require a radical transformation of the AI research enterprise, facilitated by significant and sustained investment. These are the major recommendations of a recent community effort coordinated by the Computing Community Consortium and the Association for the Advancement of Artificial Intelligence to formulate a Roadmap for AI research and development over the next two decades.

Societal Benefits of AI

AI systems have the potential for transformative impact across all sectors of society and for substantial innovation and economic growth. At the same time, there are many concerns about the security and vulnerability of systems with these capabilities, as well as about the future of work in such a world. The Roadmap articulates AI benefits in several specific areas: 1) boost health and quality of life, 2) provide lifelong education and training, 3) reinvent business innovation and competitiveness, 4) accelerate scientific discovery and technical innovation, 5) expand evidence-driven social opportunity and policy, and 6) transform national defense and security. The Roadmap includes detailed vignettes describing how AI innovations could impact individual lives, organizations, and society. Potential AI innovations include health monitors and advisors, mental and behavioral health coaches, enhanced education for remote students, effective natural disaster response, rapid materials discovery, accurate models of water resources, expeditious cross-disciplinary medical research, business innovation in personal devices, supply chain delay resolution, and resilient cyber-physical systems. All these innovations require a range of fundamental research advances in key areas of AI.

Research Priorities to Realize Societal Benefits

Major research priorities that arise from motivating societal drivers include:

- ▶ **Integrated intelligence**, including developing foundational principles for combining modular AI capabilities and skills, approaches for contextualizing general capabilities to suit specific uses, creation of open shared repositories of machine-understandable world knowledge, and understanding human intelligence both to inspire novel AI approaches and to develop models of human cognition.
- ▶ **Meaningful interaction**, comprising techniques for productive collaboration in mixed teams of humans and machines, combining diverse communication modalities (verbal, visual, emotional) while respecting privacy, responsible and trustworthy behaviors that can be corrected directly by users, and fruitful online and real-world interaction among humans and AI systems.
- ▶ **Self-aware learning**, developing robust and trustworthy learning, quantifying uncertainty and durability, learning from small amounts of data and through instruction, incorporating prior knowledge into learning, developing causal and steerable models from numerical data and observations, and learning real-time behaviors for intentional sensing and acting.

Challenges of the Current AI Landscape

Through the Roadmap activities, several critical challenges were identified. To begin, the field has matured beyond its initial academic focus on algorithms and theories and into a context of continuous data collection, social and interactive experimentation, and massive amounts of knowledge about a constantly changing world. Building from those foundations, the tech industry has compiled and leveraged massive resources—datasets, knowledge graphs, special-purpose computers, and large cadres of AI engineers—to propel powerful innovations. Tackling the research priorities above will require appropriate resources that can drive basic research of a more experimental nature. Without the right resources, academic AI research is limited—without answers to foundational questions, AI applications in industry will bring limited innovations. The constraints, incentives, and timelines are very different, too: Industry is largely driven by practical, near-term solutions, while academia is where many of the fundamental long-term questions are asked. Moreover, AI challenges span all areas of computer science and computer engineering, as well as cognitive science, psychology, biology, mathematics, public policy, ethics, education, and communication, to name just a few. The talent pool is another critical issue in the current AI ecosystem: The need for AI expertise far exceeds the supply and the gap will only continue to grow if not addressed. Many AI faculty have moved to industry to pursue new opportunities brought about by unique data and massive resources. US PhD graduates find attractive opportunities abroad as Asia and Europe are making multi-billion dollar investments in this area. Finally, there are many concerns about the security and vulnerability of AI systems, ensuring ethical uses of AI, and the future of work.

Recommendations

Surmounting these challenges will require a reinvention of the AI research enterprise to create a comprehensive national AI infrastructure and re-conceptualize AI workforce training. To that end, the Roadmap offers the following specific recommendations:

I – Create and Operate a National AI Infrastructure to serve academia, industry, and government through four interlocking capabilities:

- ▶ **Open AI platforms and resources:** a vast interlinked distributed collection of “AI-ready” resources (such as curated high-quality datasets, software, knowledge repositories, testbeds for personal assistants and robotics environments) contributed by and available to the academic research community, as well as to industry and government.
- ▶ **Sustained community-driven AI challenges:** organized sequences of challenges that build on one another, posed by AI and domain experts to drive research in key areas, building upon—and adding to—the shared resources in the Open AI Platforms and Facilities.
- ▶ **National AI Research Centers:** multi-university centers with affiliated institutions, focused on pivotal areas of long-term AI research (e.g., integrated intelligence, trust, and responsibility), with decade-long funding to support on the order of 100 faculty, 200 AI engineers, 500 students, and necessary computing infrastructure. These centers would offer rich training for students at all levels. Visiting fellows from academia, industry, and government will enable cross-cutting research and technology transition.
- ▶ **Mission-Driven AI Laboratories:** living laboratories for AI development in targeted areas of great potential for societal impact. These would be “AI-ready” facilities, designed to allow AI researchers to access unique data and expertise, such as AI-ready hospitals, AI-ready homes, or AI-ready schools. They would work closely with the National AI Research Centers to provide requirements, facilitate applied research, and transition research results. These laboratories would be crucial for R&D, dissemination, and workforce training. They would have decade-long funding to support on the order of 50 permanent AI researchers, 50 visitors from AI Research Centers, 100-200 AI engineers and technicians, and 100 domain experts and staff.

II – Re-conceptualize and Train an All-Encompassing AI Workforce, building upon the National AI Infrastructure listed above to:

- ▶ **Develop AI Curricula at All Levels:** guidelines should be developed for curricula that encourage early and ongoing interest in and understanding of AI, beginning in K-12 and extending through graduate courses and professional programs.
- ▶ **Create Recruitment and Retention Programs for Advanced AI Degrees:** including grants for talented students to obtain advanced graduate degrees, retention programs for doctoral-level researchers, and additional resources to support and enfranchise AI teaching faculty.
- ▶ **Engage Underrepresented and Underprivileged Groups:** programs to bring the best talent into the AI research effort.
- ▶ **Incentivize Emerging Interdisciplinary AI Areas:** initiatives to encourage students and the research community to work in interdisciplinary AI studies—e.g., AI safety engineering, as well as analysis of the impact of AI on society—will ensure a workforce and a research ecosystem that understands the full context for AI solutions.
- ▶ **Highlight AI Ethics and Policy:** including the importance of the area of AI ethics and policy, and the imperative of incorporating ethics and related responsibility principles as central elements in the design and operation of AI systems.
- ▶ **Address AI and the Future of Work:** these challenges are at the intersection of AI with other disciplines such as economics, public policy, and education. It is important to teach students how to think through the ethical and social implications of their work.
- ▶ **Train Highly Skilled AI Engineers and Technicians:** support and build upon the National AI Infrastructure to grow the AI pipeline through community colleges, workforce retraining programs, certificate programs, and online degrees.

III – Core Programs for basic AI Research are critical. The new resources and initiatives described in this Roadmap cannot come at the expense of existing programs for funding AI research. These core programs—which provide well-established, broad-based support for research progress, for training young researchers, for integrating AI research and education, and for nucleating novel interdisciplinary collaborations—are critical complements to the broader initiatives described in this Roadmap, and they too will require expanded support.

All of this will require substantial, sustained federal investment over the course of the 20-year period covered by this Roadmap, but the outcomes will be transformative. The recommendations above are not only a scaffold for interdisciplinary, forward-looking R&D that will drive scientific and economic advances while taking into consideration issues around security, vulnerability, policy, and ethics. The recommendations in this Roadmap will also allow the retention of the best talent in fertile research settings, creating extensive human capital in this crucial technology area—another important benefit to society and the economy.

A 20-Year Community Roadmap for Artificial Intelligence Research in the US

Volanda Gil (USC) and Bart Selman (Cornell), co-chairs

Few of AI's challenges can be solved as piecemeal academic research projects

ARTIFICIAL INTELLIGENCE (AI) LANDSCAPE

- Data-driven AI methods are highly effective but have important flaws
- Industry focuses largely on practical, near-term solutions using massive proprietary resources
- Academia asks many of the fundamental long-term questions that lay the foundations for AI

ASPIRATIONS

- Reduced healthcare cost
- Universal personalized education
- Evidence-driven social opportunity
- Accelerated scientific discovery
- Unprecedented innovation for businesses
- National defense and security

AI-driven capabilities

- Behavioral health coaches
- High payoff experiments
- Opportunistic education
- Resolve supply chain delays
- At-home robot caregivers/helpers
- Effective natural disaster response
- Novel business processes
- Address food and water insecurity
- Resilient cyber-physical systems

CROSS-CUTTING ISSUES

- Security & vulnerability, ethics, resources (data, hardware, software, storage, people...)

AI RESEARCH PRIORITIES

Integrated Intelligence

- Science of integrated intelligence
- Contextualized AI
- Open knowledge repositories
- Understanding human intelligence

Meaningful Interaction

- Collaboration
- Trust and responsibility
- Diversity of interaction channels
- Improving online interaction

Self-Aware Learning

- Robust and trustworthy learning
- Deeper learning for challenging tasks
- Integrating symbolic and numeric representations
- Learning in integrated AI/robotic systems

RECOMMENDATIONS FOR AI RESEARCH

National AI Infrastructure

Core AI Programs

Workforce Training

Interdisciplinary AI

National AI Research Centers

Open AI Platforms and Facilities

Mission-Driven AI Laboratories

Broadening AI curriculum

Community-driven challenges

Recruitment and training programs

Engaging underrepresented groups