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

Executive Summary

March 2019

Co-chairs: Yolanda Gil and Bart Selman

Decades of artificial intelligence (AI) research 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, converse about order placement, and control cars. The ubiquitous deployment of AI systems has not only created a trillion dollar AI industry that is projected to quadruple in three years, but has also exposed the need to make AI systems fair and trustworthy as well as more competent about the world in which they (and we) operate. Future AI systems will be rightfully expected to handle complex tasks and responsibilities, engage in meaningful communication, and improve their awareness through experience.

The next generation of AI systems have the potential for transformative impact on society. For example, lifelong personal assistants will enable an elderly population to live longer independently, AI health coaches will provide advice for lifestyle choices, customized AI tutors will broaden education opportunities, and AI scientific assistants will dramatically accelerate the pace of discovery. Recognizing that AI will be a major driver of the economy over the next several decades, Asia and Europe are making multi-billion dollar AI investments. With strategic investments, the US can remain the preeminent leader in AI and benefit from its broad societal and economic impacts.

Achieving the full potential of AI technologies poses research challenges that will require significant sustained investment and a radical transformation of the AI research enterprise. This is the main finding of a recent study by leading AI experts carried out 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.

The 20-year research roadmap for AI envisions three major areas of significant potential impact:

- *Integrated intelligence*, including foundations for principled combination of modular skills and capabilities, contextualization of general capabilities to suit particular uses, and creation and use of open shared repositories of machine understandable world knowledge.
- *Meaningful interaction*, comprising productive collaboration, diverse communication modalities, responsible and trustworthy behaviors, and fruitful online and real-world interaction.
- *Self-aware learning*, ranging from robust and trustworthy learning, learning from few examples and through instruction, developing causal and steerable models from numerical data and observations, and real-time intentional sensing and acting.

Underlying these research directions is the quest to understand intelligence in all forms (artificial, human, animal) and contexts. The AI community is eager to pursue this research agenda, but there are major impediments to making substantial headway.

First, the field of AI has reached a maturity level that goes beyond the initial academic focus on algorithms and theories and into embracing live instrumented deployments, continuous data collection, social and interactive experimentation, dynamic environments, and massive amounts of knowledge about a constantly changing world. This requires new facilities that do not exist in academia today. Although major AI innovations have roots in academic research, universities now lack the massive resources (unique datasets, special-purpose computing, extensive knowledge graphs, well-trained AI engineers, etc.) that have been acquired or developed by major IT companies. These are fundamental capabilities to build forward-looking AI research programs. This also puts universities at a serious disadvantage in terms of attracting talented graduate students and retaining influential senior faculty. Moreover, because AI resources in major IT industry labs are generally proprietary, this uneven playing field also negatively affects smaller businesses and non-IT industry sectors, as well as government organizations, all of which have traditionally benefitted from the open nature of academic research.

Second, research requires highly interdisciplinary teams that can only succeed in long-term sustained programs that are currently rarely available. AI challenges span all areas of computer science, as well as cognitive science,

psychology, biology, mathematics, social and natural sciences, engineering safety, public policy, ethics, education, and communication, to name a few.

Third, the overwhelming demand and excitement surrounding data-rich machine learning has caused a redirection of faculty, students, and curricula towards this particular area of AI, at the expense of other AI areas that are crucial for a broader research roadmap.

Fourth, the need for AI expertise surpasses current production of university graduates with AI skills at the undergraduate, masters, and PhD levels. Many PhD level AI graduates in the US find attractive opportunities abroad. Although there is great demand for AI practitioners and data scientists, universities lack the facilities to prepare students properly for industry settings.

Achieving this vision will require a reinvention of the AI research enterprise to create a comprehensive national AI infrastructure and to re-conceptualize AI workforce training. The 20-year AI research roadmap includes the following specific recommendations:

I --- Create and operate a National AI infrastructure to serve academia, industry, and government through four interlocking capabilities:

a) *Open AI Platforms and Facilities*: a vast interlinked distributed collection of "AI-ready" resources (growing datasets, software libraries, knowledge repositories, instrumented homes and hospitals, robotics environments, cloud services, etc.) contributed by and available to the research community as well as to industry and government.

b) *Sustained Community-Driven AI Challenges*: organizational structures that coordinate the formulation by researchers of well-defined challenge roadmaps to jointly address key problems, reformulate them in unison with new advances, promote integration and well-engineered systems, and create shared resources in the Open AI Platforms and Facilities.

c) *National AI Research Centers*: concentrations of first-class researchers, including multi-year funded Faculty Fellows affiliated with a range of academic institutions and Industry Resident Fellow positions from other organizations. Together, these will create critical mass to address core AI research challenges and facilitate technology transfer to industry.

d) *National AI Laboratories*: government organizations that will provide sustained capabilities and AI experts to support the Open AI Platform and AI Challenges, and address vertical sectors of public interest such as health, policy, education, and science.

II --- Re-conceptualize and train an all-encompassing AI workforce

a) *Recruitment programs for AI,* including grants for talented students to obtain advanced graduate degrees, retention programs for doctorate-level researchers, and engagement of underrepresented groups.

b) *Broadening AI curriculum and incentivizing non-traditional and interdisciplinary AI studies*, with priority to AI policy and law, as well as AI safety engineering.

c) *Training highly skilled AI engineers* and data scientists through the Open AI Platform, thereby significantly growing the pipeline through community colleges and workforce retraining programs.

The combination of shared infrastructure capabilities and a massive skilled workforce will put the United States in a unique position to continue to be the world's leader in AI research, development, and technology transfer, by 1) pursuing the forward-looking AI research that will lead to sustained and broad innovations in AI; 2) creating unique, comprehensive, and effective AI capabilities; 3) attracting and retaining the best talent in fertile research settings; 4) creating extensive human capital in this crucial technology area; and 5) driving AI technologies to address important problems in sectors less favored by industry, such as scientific discovery, education, and public policy.

This is the result of the community road mapping effort which took place in the fall 2018 and winter of 2019. Information about the AI road mapping effort can be found here (https://cra.org/ccc/visioning/visioning-activities/2018-activities/artificial-intelligence-roadmap/). This material is based upon work supported by the National Science Foundation under Grant No. 1136993. *Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.*

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

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

AI LANDSCAPE

- Exploiting data with AI/ML highly effective but limited use
- Industry push of AI mostly for consumer products
- Piecemeal funding programs and academic projects
- IT giants amassing significant proprietary resources and experts
- Rest of industry, government, academic lack access to large AI infrastructure and engineering resources

ASPIRATIONS

AI-driven capabilities:

- Mental and behavioral health coaches
- Accurate models of water resources
- Speed up materials science experiments
- Augment education for remote students
- Resolve supply chain delays
- At-home robot caregivers/helpers
- Response for natural disasters
- Collaborative omics discoveries
- Train for robot repair jobs
- Businesses innovation in personal devices
- Game design startups
- Scientific models from theories and data
- Improve law enforcement and training
- Resolve food insecurity and distribution
- Resilient cyber-physical systems

- Reduce healthcare cost
- Accelerating scientific discovery
- Universal personalized education
- Unprecedented innovation for businesses
- Evidence-driven social opportunity

AI challenges not solvable as piecemeal academic research projects

AI challenges not a priority for big AI/IT industry players

AI RESEARCH PRIORITIES

Integrated intelligence

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

Meaningful interactions

- Collaboration
- Trust and responsibility
- Diversity of interaction channels
- Improved interaction between people

Self-aware learning

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

RECOMMENDATIONS

Audacious AI Research

National AI Infrastructure

- National AI Labs
- National AI Research Centers
- Open AI Platforms and Facilities
- Sustained community-driven challenges

Workforce Training

- Interdisciplinary AI
- AI faculty and student retention
- Broadening AI research priorities
- AI engineers workforce (retraining, CC, MS)

9 March 2019