The Defense Advanced Research Projects Agency's Artificial Intelligence Vision

Peter Highnam

> "It seems reasonable to envision… bringing computing machines effectively into processes of thinking that must go on in real time, time that moves too fast to permit using computers in conventional ways. To think in interaction with a computer in the same way that you think with a colleague whose competence supplements your own, will require much tighter coupling between man and machine than is possible today."
>
> *– J.C.R. Licklider (1960), First Director of the ARPA Information Processing Techniques Office*

Licklider's vision set in motion six decades of broad multidisciplinary research. His foresight is quintessential DARPA (Defense Advanced Research Projects Agency) — he saw what was both technically plausible and what was necessary for the US Department of Defense, and he inspired generations of researchers to execute his vision. DARPA's sustained investments have involved hundreds of DARPA program managers and many thousands of engineers and scientists working with the agency. Not all

Copyright © 2020, Association for the Advancement of Artificial Intelligence. All rights reserved. ISSN 0738-4602

Knowledge Technologies at DARPA: The Next 20 Years

While writing "A 20-Year Roadmap for Artificial Intelligence Research in the US," many of our discussions centered on reshaping the artificial intelligence (AI) research landscape and the infrastructure available to tackle larger challenges than ever before. I had a chance to reminisce about how DARPA has left a mark on entire areas of AI, in particular my own research area of knowledge technologies.

I watched DARPA influence the planning community through the ARPA-Rome Lab Planning Initiative. When I joined the program, I had just finished my PhD thesis at Carnegie Mellon University on learning planning knowledge by experimentation. In an era where the research community was focused on automated planning and learning algorithms, my first DARPA meeting immediately opened my eyes to the wider world of human-guided planning, learning planning knowledge from users, and knowledge-rich decision-making. My own research focused on plan critiquing and comparison, and the importance of AI systems that have substantial knowledge of planning practices, interacting goals, and situation context to give users thorough feedback about potential courses of action. I remember vividly a staged simulation for an evacuation off a fictitious island that highlighted the need for distributed agent coordination and constraint negotiation. DARPA helped the research community understand and prioritize many novel and important research areas in planning, with long-lasting impact. These research areas are still full of unsolved problems, as if waiting for imaginative DARPA programs to push them forward.

Three DARPA programs shaped knowledge technologies in the 1990s and early 2000s. The High-Performance Knowledge Bases and Rapid Knowledge Formation programs led to significant methodologies for knowledge capture, and pioneered the idea of starting with basic textbook knowledge and passing high-school and college tests as a means to evaluate the amount and quality of knowledge systems. Through continuous evaluations of question answering and problem-solving capabilities, these programs exposed the benefits of interconnected assertions and knowledge units, the need for users to understand and trust the knowledge sources used for reasoning, and the challenges of supporting continuous updates and extensions. My research concerned supporting users to update knowledge by ensuring the changes were consistent and complete. Soon after, the DARPA Agent Markup Language program then drove the community to develop standards for knowledge sharing, brokering intense joint work with European initiatives in this area and eventually leading to World Wide Web recommendations that have shaped the Semantic Web and linked data. These standards have profoundly changed scientific research (with numerous community-developed ontologies in many disciplines) and have become a key technology in industry (known as knowledge graphs). The recent articulation of open knowledge networks will push this area forward significantly, and DARPA will have opportunities to continue to play a unique role.

In the 20-year AI research roadmap, there are recommendations for national AI centers — both large multiuniversity centers for important research themes and mission-driven living laboratories for different application areas. The Personal Assistants that Learn program provides a unique example of how significant multi-institutional projects can be set up in AI. The Cognitive Agent that Learns and Organizes project had hundreds of researchers from two dozen universities working toward the common goal of interactive support for daily tasks. My own group focused on assistance with to-do lists, and we guided volunteers to provide thousands of assertions to capture common knowledge about the world that could be used to automatically reorganize and manage to-dos. That basic research on "volunteer knowledge collection," which pioneered what is now known as crowdsourcing, is an example of how DARPA programs support basic research agendas driven by hard practical problems. A unique feat of the Cognitive Agent that Learns and Organizes was the broad scope and diversity of basic AI research that was unified under a single project with common goals.

DARPA has long recognized the unique role of its programs to shape entire areas of AI. As I reflect on the recommendations of our roadmap for the next 20 years, AI needs substantial multiuniversity multidecade-funded centers to tackle some of the wide-ranging challenges in our field. The Cognitive Agent that Learns and Organizes project is arguably the closest example that I have seen to this, because its 5-year $150M investment dwarfs the largest National Science Foundation and Multidisciplinary University Initiative awards in AI. While it had a strict evaluation tempo handled by a focused team, it also allowed an entire ecosystem of basic research to flourish. I am hopeful that DARPA will see a role in funding large AI projects that blend basic, experimental, and practical research.

– Yolanda Gil

the investments succeeded, but the new science and technologies they produced have changed the world.

It has not been a smooth ride. There were periods of advances and great optimism with attendant promises and expectations, and periods of disappointment and slow progress. As pointed out by Herbert Simon in 1969 (Simon 1969), the field's name has significant baggage for technical and lay persons. This baggage provides impediment and inspiration. DARPA's support for the field has persisted throughout, albeit with significant fluctuations.

Over the last few years DARPA has adopted a perspective on the field's development that is descriptive and useful. The agency's "three waves of AI technologies" neatly tags the dominant developments, which correlate well with the availability of computation and communications. The waves are covered in depth elsewhere in this issue; it suffices to observe here that AI technologies, general computation, and computer communications advance together. As we peer into the near future, third-wave AI contextual reasoning surely drives and will be enabled by fifth-generation ("5G") communication technologies and what follows that, human-centric computing devices, and widespread stratified cloud-based compute and knowledge.

Heeding the earlier observation of the perils of overpromising, there remains much to do to make today's AI technologies robust and dependable enough to be widely integrated into massive enterprise systems and to be used routinely in safety-critical contexts. Some of the work is deeply technical, such as handling pathologic failure modes, mitigating bias, and defense against adversarial attack. As much work, if not more, remains to establish the tradecraft and tools of system engineering when emerging AI technologies are inserted into enterprise systems.

DARPA continues to drive the development of AI technologies. At the same time, DARPA is incorporating new technologies into military systems that enable collaboration with warfighters, resulting in better decisions in complex, time-critical, battlefield environments; allow shared understanding of massive, incomplete, and contradictory information; and empower unmanned systems to perform critical missions safely and with high degrees of autonomy. In short, we're still tracking with Licklider's vision.

In September 2018 DARPA announced AI Next, a multiyear investment of more than $2 billion. Key campaign areas include developing robust foundations for second-wave technologies, aggressively inserting second-wave AI technologies into systems, and exploring and creating third-wave AI science and technologies. Human–machine symbiosis is not out of reach, but making it a reality for the US Department of Defense and beyond will require steady, targeted R&D investments. DARPA plans to continue driving that investment well into the future.

Reference

Simon, H. 1969. *The Sciences of the Artificial*. Cambridge, MA: The MIT Press.

Peter Highnam is deputy director of DARPA. Highnam helps advance critical research in several fields including AI, data analytics, communications, high performance computing, electronic warfare, biosecurity, and warfighter health. Highnam previously worked in the Office of the Director of National Intelligence, serving as Research director at the National Geospatial-Intelligence Agency and in the Intelligence Advanced Research Projects Activity as office director and then director. He was a senior advisor at the National Institutes of Health and the Biomedical Advanced Research and Development Authority. After studying computer science at Carnegie Mellon University he worked for eleven years for Schlumberger Limited, before joining DARPA (the first time) as an electronic warfare program manager.