DARPA's Impact on Artificial Intelligence

Scott Fouse, Stephen Cross, Zachary J. Lapin

■ *The Defense Advanced Research Project Agency's (DARPA) mission is to make pivotal investments leading to research breakthroughs that support national security. DARPA artificial intelligence (AI) programs have emphasized the need for machines to perceive and interact with the world around them; to frame problems and to arrive at solutions and decisions based on reasoning; to implement those decisions, perhaps through consultation with a human or another machine; to learn; to explain the rationale for decisions; to adhere to rules of ethical behavior defined for humans; to adapt to dynamic environments; and, to do all of this in real-time. In short, DARPA has always been interested in AI frameworks that integrate AI and computer science technologies, and the application of those frameworks to DARPA-hard problems. In this article, we describe the significant role that DARPA has played in the establishment of AI, and introduce six articles that explore DARPA's Three Waves of AI.*

In 1962, J.C.R. Licklider created the US Information Processing Techniques Office at the Advanced Research Projects Agency (ARPA). His vision, published two years earlier in his seminal work *Man–Computer Symbiosis* (Licklider 1960), heralded an ambitious, and ultimately successful, push to develop artificial intelligence (AI) technologies. The Agency, now called DARPA with the *D* emphasizing its focus on defense applications, has supported AI research, as popularity has ebbed and flowed, over the past 60 years.

DARPA's job is to change what's possible — to do the fundamental research, the proof of principle, and the early stages of technology development that take impossible ideas to the point of implausible and then, surprisingly, possible. No other US Department of Defense agency has the mission of working on projects with such a high possibility of producing truly revolutionary new capabilities — or such a high possibility of failure. A big part of DARPA's expertise is seeking high-payoff capabilities by managing risk in ways that help keep the innovation pipeline flowing. These are the kinds of foundational-research efforts that promise to impact national security like the ARPAnet, which yielded the internet, and AI research, which promises to have transformative effects for decades to come.

Over its six-decade history, DARPA has made significant investments in the creation of AI technologies that have produced game-changing capabilities for the US Department of Defense and beyond. In September 2018, I unveiled a multiyear investment of more than $2 billion DARPA would make in new and existing programs called the AI Next campaign. This commitment will define scientific and technical exploration, as well as resulting military capabilities, for the next century.

In addition to new and existing DARPA research, a key component of the AI Next campaign is DARPA's Artificial Intelligence Exploration program. Artificial Intelligence Exploration constitutes a series of high-risk, high-payoff projects where researchers will work to establish the feasibility of new AI concepts within 18 months of award. Leveraging streamlined contracting procedures and funding mechanisms is enabling these efforts to move from proposal to project kickoff within three months of an opportunity announcement.

AI Next and Artificial Intelligence Exploration are substantive efforts to address the limitations of existing AI. DARPA believes current research and development investments around the world are much too focused on second-wave AI or machine learning, which is particularly good in finding patterns in voice and images and has many commercial applications. In the United States, however, we are pursuing programs that will make second-wave AI more robust for defense and security applications, all while helping realize the third wave of AI or contextual reasoning. Incorporating these technologies in future systems promises to facilitate better decision-making; enable shared understanding of massive, incomplete, and contradictory information; and empower unmanned systems to perform critical missions safely.

The United States enjoys a major competitive advantage in its access to the world's best universities, government laboratories, and industry partners. Marshalling this unique mix of resources, we are well poised to be the first to forge new AI theories and methods that will make it possible for machines to adapt contextually to changing situations, advancing computers from mere tools to true collaborative partners. Going forward, DARPA will be unafraid about exploring these new capabilities — DARPA's core mission — pushing scientific frontiers ahead of adversaries.

– Steven H. Walker

AI is currently *in*, and it is definitely big business, with current estimates suggesting that "… AI will add $13 trillion to the global economy over the next decade." (Fountaine and Saleh 2019) DARPA is now investing more than $2 billion in the next generation of AI through its initiative, *AI Next*. Advances will not only require simultaneous and coordinated progress in knowledge representation and reasoning, machine learning, human language technology, and vision and robotics, but also in the tight integration of each component to realize trustworthy intelligent systems able to operate autonomously or team with humans. The realization of AI Next rests on the AI community both to innovate as performers and to serve at DARPA; over the next five years, AI Next will not only push the limits of current technology, but also challenge human creativity to reimagine what AI can, and should, do. Former DARPA Director, Steven Walker, shares his vision for AI Next in an accompanying sidebar.

While DARPA research programs are surging forward, in this collection we are looking back at the significant role that the Agency played to establish AI. DARPA's mission is to make pivotal investments leading to research breakthroughs that support national security. DARPA AI programs have emphasized the need for machines to perceive and interact with the world around them; to frame problems and to arrive at solutions and decisions based on

reasoning; to implement those decisions, perhaps through consultation with a human or another machine; to learn; to explain the rationale for decisions; to adhere to rules of ethical behavior defined for humans; to adapt to dynamic environments; and, to do all of this in real-time. In short, DARPA has always been interested in AI frameworks that integrate AI and computer science technologies, and the application of those frameworks to DARPA-hard problems.

DARPA's support for AI research can be visualized by considering significant, broad, research focus areas — DARPA's Three Waves of AI. The first wave is the era of handcrafted, declarative knowledge; the second wave is statistical learning; and the third wave is the future. DARPA termed the third wave *Contextual Adaptation*, but we are hesitant to label the future and instead refer to it by DARPA's initiative, AI Next. In addition to the three waves, we have identified six major phases of DARPA's AI investment: AI Beginnings, Strategic Computing, Knowledge/Planning, Cognitive Systems, Data Analytics, and AI Next. The waves, the phases, and a selection of historically significant programs are illustrated in figure 1, which serves as the foundation for all articles in this collection.

These articles form a historical perspective of DARPA's influence on AI over six decades. As a result, the articles focus more on the past than the present. They do not include more recent and current AI programs that promise to continue DARPA's outsized influence on AI and form the basis for the third wave, including Big Mechanism, Communicating with Computers, Probabilistic Programming for Advanced Machine Learning, Data-Driven Discovery of Models, Learning with Less Labels, Explainable Artificial Intelligence, Competency-Aware Machine Learning, Machine Common Sense, Life-Long Learning Machines, Guaranteeing AI Robustness against Deception, Science of AI and Learning for Open-World Novelty, Media Forensics, Semantic Forensics, Radio Frequency Machine Learning Systems, Machine Reading, Broad Operational Language Translation, Robust Automatic Transcription of Speech, Low Resource Languages for Emergent Incidents, Active Interpretation of Disparate Alternatives, Knowledge-directed Artificial Intelligence Reasoning Over Schemas, Grounded Artificial Intelligence Language Acquisition, and other smaller AI Exploration topics.

The holistic approach to integrating subfields of AI, taken by DARPA, and advocated by the authors of this collection, has also been a consistent theme in Association for the Advancement of Artificial Intelligence presidential addresses. Allen Newell cautioned against a fragmented view when he stated in the first Association for the Advancement of Artificial Intelligence presidential address that:

> The first watchword is cooperation — cooperation with all the others that are trying to nurture artificial intelligence. (Newell 1980)

Raj Reddy also challenged us to embrace futuristic grand challenges that require simultaneous advances across subfields of AI as well as in computer science in his 1988 Presidential Address (Reddy 1988):

> … success in AI depends on advances in all of computer science. We are not, and never have been an island unto ourselves. Finally, all parts of AI belong together. Success in AI requires advances in all of its disparate parts including chess, cognitive science, logic, and connectionism.

Similarly, Barbara Grosz advocated for continued progress in collaborative systems (Grosz 1994), Ron Brachman discussed the challenges of achieving intelligent behavior in integrated systems (Brachman 2006), and Tom Dietterich challenged the AI community to achieve truly intelligent and robust systems (Dietterich 2017). Another recent Association for the Advancement of Artificial Intelligence President, Manuela Veloso, stressed the need for more research into symbiotic autonomy as another view on the integration theme (Brandom 2016).

What can you do? We believe the time has come for each of us to become a responsible spokesman for the entire field. This requires some additional effort on our part to be articulate and be convincing about the progress and prospects of AI. Finally, choose your favorite grand challenge relevant to the nation, and work on it.

The significant AI successes of DARPA, presented throughout this collection, are attributable to the unique ways in which DARPA supports research. DARPA brings together communities of critical mass to attack and solve hard problems by motivating system-level grand challenge visions of national significance. It employs test beds, competitions, and other means to bring communities together and to accelerate progress. A former DARPA program manager, office director, and Agency director — Arati Prabhakar — shares her views of the importance of DARPA in an accompanying sidebar.

DARPA successes can be largely attributed to the program managers. DARPA hires exceptional program managers, and gives them the freedom and resources to exploit their passion. Program managers come from diverse civilian and military backgrounds, and all have demonstrated technical and operational excellence. Motivation is not only internal to the program managers, but also amplified by DARPA's aggressive timeline: Program managers have an average of only four years to create their impact. Bob Engelmore (Engelmore 1981) was the first editor of *AI Magazine* and his 1982 article about his time as a program manager is still very relevant. Many members of the AI research community have served two or more tours at DARPA (for example, Randy Garret and Dave Gunning served three times; Allen Sears, Howie Shrobe, Ted Senator, and Charles Wayne have each served as program managers twice).

Program managers actively engage scientific and technical communities, and energize them with

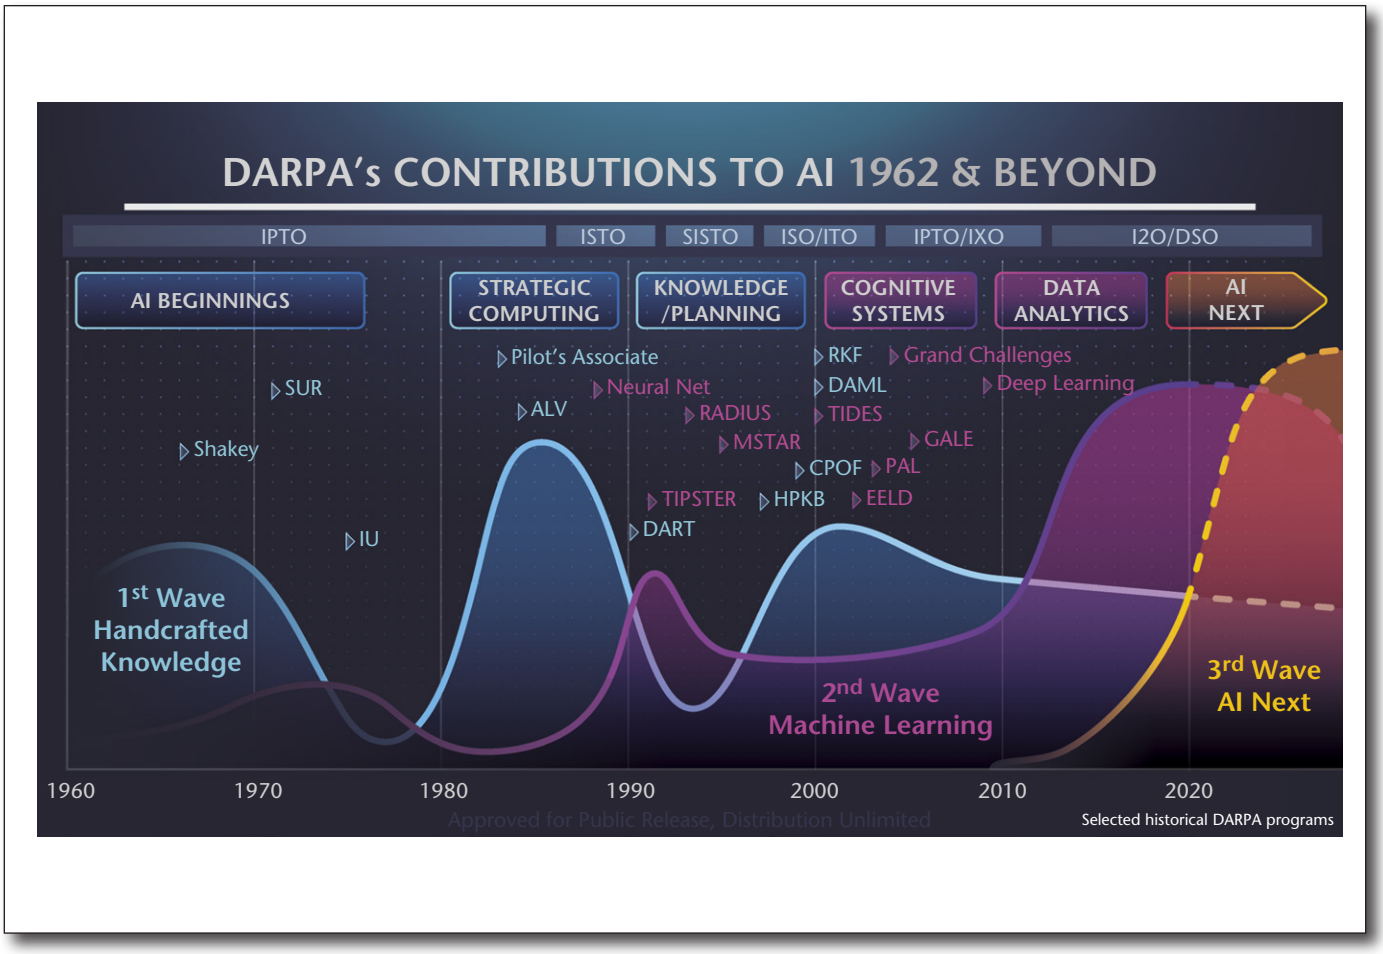

Figure 1. DARPA's Contributions to AI.

A graphical representation of DARPA's contribution to AI beginning in 1962. DARPA's Three Waves are shown in blue, purple, and yellow, with the height of each wave representing DARPA's investment. Selected historical DARPA AI programs discussed in this collection are shown. (*Figure courtesy of DARPA.*)

exciting challenges and opportunities: ones for which success is uncertain and perhaps not yet even fully defined. DARPA's drive is for innovation, and while community members are grounded in deep scientific principles, they cannot allow their knowledge and experience to hold them back — community members study user needs to generate novel concepts for game-changing capabilities. These communities must also have broad system experiences, maximizing the potential for the integration of developed technologies with existing systems.

DARPA is focused on problems, not requirements. This is a subtle but especially important distinction, and highlighted by DARPA's development and use of Broad Agency Announcements rather than the traditional Request For Proposal. The requirements of Requests For Proposals are tied to current processes and metrics, and often constrain innovation to evolutionary change. Broad Agency Announcements, on the other hand, facilitate the discovery of new processes, which yield new metrics and lead to revolutionary change.

This collection is devoted to DARPA's impact on AI over the past sixty years. As suggested, a consistent theme has been the integration and scaleup of AI subfields. Many AI successes now in daily use resulted from breakthroughs produced by pioneering DARPA programs driven by the far-reaching system visions articulated by Licklider and those who followed him.

This collection consists of six articles. The first five are authored by a former DARPA official and a well-known principal investigator: Richard Fikes and Tom Garvey discuss progress in knowledge representation and reasoning. Josh Alspector and Tom Dietterich review the progress in machine learning. Charles Wayne and Mark Liberman review the progress in human language technology. Rama Chellappa, Tom Strat, and Vishal Patel discuss progress in robotics and image understanding. Ron Brachman, David

During my time as director of DARPA — 2012 to 2017 — artificial intelligence was starting its renaissance as years of machine learning research enabled a raft of new solutions. It was the classic overnight success, after decades of hard work. But something was different this time: What emerged was a generation of technology that interacts intimately with humans, in a way that wasn't so for the age of semiconductors and computers and networks. We now have machines that understand and act on social media posts, machines that see photos as collections of faces rather than collections of pixels, machines that a child can converse with, or that can approve loans or set bail or drive cars. They're in our human sphere now, and it's up to us to make the choices that steer away from dystopia and toward Licklider's optimistic vision of human-machine symbiosis.

At DARPA, we developed a narrative about AI's coming of age, Three Waves of AI, to explain to ourselves and to others what was unfolding — both the power and the limitations of the machine learning wave. I used it in my many, many meetings with senior US DoD leaders and general officers as they wrestled with the military and national security implications of these advances. When I look back at those discussions, I'm struck by two observations. One is this: I came to realize that, contrary to the general perception of a Strangelovian or Terminator mindset, our military officers are trained and deeply grounded in an ethic that puts the responsibility for life-and-death battlefield decisions squarely on their own shoulders. It's a burden they bear on behalf of each of us, and one they have no intention of surrendering to machines.

That's the good news. I also saw another more subtle and, to me, more concerning dimension. Those who are responsible for our security are daily steeped in an environment of risk. The intelligence analyst looking for a terrorist, the soldier clearing a neighborhood, the cop on a beat, a national security strategist considering the actions of other nation states — to protect us, each must see the world through the lens of danger. A different but parallel pattern shapes behavior in companies, where the lens is the pursuit of delightfully jaw-dropping customer experiences, profits, and growth. In any of these highly focused roles, it can be nearly impossible to step back and fully consider the societal costs of successfully achieving these worthy goals.

In one obvious example, our machine-enhanced ability to observe, understand, and influence people in unprecedented ways today is posing a serious threat to individual privacy — a most American ideal that's about far more than just being left alone. As James Bennet keenly observed, "Privacy sustains space for free thought and expression, for the growth that comes from mistakes without public shame." For the foreseeable future, we are sure to be grappling with the challenge of achieving the enormous benefits of AI for security and prosperity while sustaining our values.

My personal focus today is bringing new forms of innovation to society's complex challenges for the decades ahead. The meetings are different than they were in the Pentagon; now I'm working with people who understand the oppressive cost of US healthcare, the daunting challenges so many face in reaching for economic opportunity, and the looming consequences of global climate instability. At the heart of many of these challenges are structural and systemic issues that I know from my recent DARPA experiences can start to be deconstructed and addressed using the astounding new tools at our disposal today. New types of data, new analytical and experimental techniques, and new ways to model and reason about complexity can be combined with theory and domain expertise to form hypotheses, design experiments, and ultimately to deeply understand and influence complex systems from biology to weather to social dynamics. These advances create the opportunity to unravel complex problems that have long been considered intractable.

That's what the future of AI can and should look like. The burden is on us — on our very human shoulders — to use this powerful technology to reinforce our core values, and to lift our society into a better future.

– Arati Prabhakar

Gunning, and Murray Burke survey work in integrated intelligent systems with a focus both on autonomous and computer-aided systems. Finally, this collection concludes with a vision from the Acting Director of DARPA, Peter Highnam, in which he discusses the plans for AI Next.

We also acknowledge our sincere appreciation to Bill Mark, Karen Myers, Ray Perrault, Larry Davis,

and Ralph Weischel, for conducting external reviews of the articles. We also point out Bill's longstanding service to DARPA as a principal investigator and as a member and past chair of the DARPA Information Science and Technology panel. Although space does not allow discussion of the role of Information Science and Technology, it provides another valuable way for members of the AI community to interact with DARPA Program Managers through studies supported each year.

The articles in this collection, however, cannot alone fully express what DARPA is, and why DARPA is unique. To gain a richer understanding of the DARPA Way, each article includes sidebars written by people intimately involved in DARPA AI programs. Individually, they all offer a unique perspective, and taken altogether they paint a picture of the significant role DARPA has played in advancing AI over the past 60 years.

Acknowledgments

This projected was supported by DARPA.

References

Brachman, R. 2006. (AA)AI — More Than the Sum of Its Parts. *AI Magazine* 27(4): 19-34. doi.org/10.1609/aimag.v9i4.950.

Brandom, R. 2016. "Humanity and AI Will Be Inseparable — Manuela Veloso, Head of Machine Learning, Carnegie Mellon University." The Verge, November 15, 2016. www.theverge.com/a/verge-2021/humanity-and-ai-will-be-inseparable.

Dietterich, T. 2017. Steps Toward Robust Artificial Intelligence. *AI Magazine* 38(3): 3-24. doi.org/10.1609/aimag.v38i3.2756.

Engelmore, R. 1981. Reflections on the ARPA Experience. *AI Magazine* 3(1): 11. doi.org/10.1609/aimag.v3i1.360.

Fountaine, T., and Saleh, T. 2019. Building the AI-Powered Organization. *Harvard Business Review* 97(4): 63-73. hbr.org/2019/07/building-the-ai-powered-organization

Grosz, B. 1994. AAAI-94 Presidential Address: Collaborative Systems. *AI Magazine* 17(2): 76-85. 10.1609/aimag.v17i2.1223.

Licklider., J 1960. Man–Computer Symbiosis. *IRE Transaction on Human Factors in Electronics*. HFE-1(1): 4-11. doi.org/10.1109/THFE2.1960.4503259.

Newell., A 1980. AAAI President's Message. *AI Magazine* 1(1): 1-4. doi.org/10.1609/aimag.v1i1.84.

Reddy, R. 1988. Foundations and Grand Challenges of Artificial Intelligence: AAAI Presidential Address. *AI Magazine* 9(4): 9-21. doi.org/10.1609/aimag.v9i4.950.

Scott Fouse is an independent technology consultant and former Vice President of Research for Lockheed Martin. He has worked on many DARPA AI projects from the late 1970s through 2009. He played key roles in two significant transition projects — the Dynamic Analysis and Replanning Tool, and the Command Post of the Future — which were recognized with three DARPA Contractor-of-the-Year awards in 1991, 2004, and 2005. He received an MSEE from the University of Southern California in 1978, and a BS in Physics from the University of Central Florida in 1976.

Stephen Cross is a retired Executive Vice President for Research at the Georgia Institute of Technology. He served as a DARPA program manager from 1989 to 1994 and later chaired DARPA's Information Science and Technology panel. He received his PhD with a Focus on AI from the University of Illinois at Urbana-Champaign in 1983.

Zachary J. Lapin is an associate at Booz Allen Hamilton. He currently provides technical support to a diverse range of efforts at DARPA. He received his Doctor of Sciences from the Eidgenössische Technische Hochschule Zürich for his work in nanophotonics and plasmonics in 2016.