


DARPA

DEEP LEARNING

$$n(B \cup C) = n(B) + n(C) - n(B \cap C)$$

He = 4.002602
Na = 22.989769
Ar = 39.948



Advancing National Security through Fundamental Research

DEEP LEARNING

THE NEED AND OPPORTUNITY

Gathering intelligence and insights from the vast amount of sensor data, communications traffic, and other forms of information that are generated continuously by today’s defense systems is critically important to supporting our military capabilities and operations. Warfighters and commanders alike utilize these insights to inform their decisions and devise strategies. However, the explosive growth in information generated by a proliferation of devices and sensors promises to overwhelm the highly trained human operators that are traditionally relied on to analyze, understand, and effectively act upon the data. To improve this process, analysts need help from machines – not just as simple collators of facts or crunchers of numbers, but as intelligent partners that can make sense of speech, text, images, video, and other complex data streams.

Since the 1950s, advances in electronics and computer technology have fueled an ongoing quest to replicate or at least emulate human intelligence in machines, a field commonly referred to as Artificial Intelligence (AI). DARPA recognized early in its institutional existence that AI could address a range of national security needs. To further that cause, the agency has for a half-century initiated a diversity of AI programs to combine advances in computer



science, mathematics, probability theory, statistics, and cognitive science with the goal of automating a range of capabilities usually associated with intelligence.

Although numerous techniques and technologies have emerged from DARPA’s AI programs, one outgrowth, Machine Learning (ML), has shown particular promise. Machine Learning relies on techniques that allow systems to learn from data, identify patterns, and make decisions with minimal human intervention. Traditional ML systems rely on “shallow” architectures in which inputs are processed by a single “hidden layer” before passing on an output to the end user. Instead of working autonomously, shallow-learning machines require expert human knowledge, extensive supervised training on labeled data (or data with

annotations that designate what it is), and massive computational power. Due to these limitations, shallow-learning machines typically offer only moderate performance advantages while still requiring significant human intervention.

THE DARPA SOLUTION

To help address the analyst’s quandary of too much data to make sense of, DARPA initiated the Deep Learning program in 2010. The goal was to develop “deep” machine-learning algorithms that could ingest a large unlabeled dataset, and then automatically detect and classify objects and activities using the unlabeled dataset as a teaching tool for the algorithm. Deep Learning machines would use multiple “hidden layers” to produce richer representations of the


images, text, and sounds that they received as inputs.

The fundamental hypothesis of Deep Learning is that this additional structure would allow the machine to more accurately classify inputs as a result of having identified abstract patterns in, for example, objects in an image. These abstract representations would render the objects recognizable, even when their orientation differed from anything that the machine had seen before. Instead of relying on the machine’s ability to memorize what an object looks like based on hundreds of thousands of labeled examples, machines with Deep Learning capabilities would be able to generalize and make inferences about the object. Such machines would learn with far fewer examples and emulate human-like object recognition more closely. From a defense perspective, the Deep Learning program focused on creating deeply layered machine-learning techniques for use in military applications, particularly automated or machine-assisted surveillance and text understanding.

THE IMPACT

As a basic research program, many of Deep Learning’s results and accomplishments were published in journal and conference papers that have since had tremendous and pervasive impacts on the theory and practice of AI/ML. Among the noteworthy technical accomplishments from the program’s research are:

- The development of a new statistical model for learning features in images without human supervision. A feature might be an object within the image, such as a vehicle or building. The model allows a computer to learn and identify the data structures that map onto an object autonomously. This model was shown to be superior to the existing state-of-the-art for small object images, which



In this artist’s conception, a driverless car’s sensors scan the environs to relentlessly inform the vehicle’s machine-learning (ML) system, which uses the data to guide current driving actions but also to modify its own programming and database so that the car becomes an ever safer means of transportation that can handle ever more real-world situations.

could be encountered in a low-resolution photograph.


- The development of a neural-network algorithm capable of analyzing both natural-scene images as well as natural-language sentences. The algorithm was described in a paper that won a Distinguished Application Award at the 2011 International Conference on Machine Learning.
- The development of new algorithms for capturing relationships between words, which can be used to distinguish between multiple meanings of the same word, an important capability for automated and accurate speech recognition.
- The development of innovative, deeply recursive networks that were successfully applied to both object recognition and difficult problems in text understanding, such as paraphrase detection, sentiment analysis, and the classification of semantic relationships (such as cause-and-effect and topic-and-message).

LOOKING AHEAD

Many of the Principal Investigators (PIs) that worked on the Deep Learning program have since been recognized as

today’s leaders in the field of AI. These industry-leading minds supported the program while working at academic institutions including Stanford University, New York University, Université de Montreal, Pennsylvania State University, and University of California, Berkeley.

The Deep Learning initiative, which continues in a growing portfolio of programs, produced many of the techniques and applications that are at the foundation of the rise in AI that we are witnessing today. Deep Learning technology is being operationalized by all major information technology companies (Google, Microsoft, Facebook, etc.) as the basis for new services and efficiencies, contributing significantly to their work with image recognition, speech translation, robotics, and facial recognition technologies. In addition, Deep Learning technologies have entered the supply chain for military systems, and will contribute to the capabilities developed through the Department of Defense’s Joint AI Center (JAIC) and other DOD initiatives.



The Machine Common Sense (MCS) program, part of a more general Deep Learning portfolio, seeks to address the challenge of articulating and encoding human common sense reasoning for intelligent machines. The program’s goal is to create deep-learning capabilities that can transcend limitations of the narrowly focused AI applications of today thereby enabling more versatile AI applications.



www.darpa.mil