Integrated Artificial Intelligence Systems

Ron Brachman, Dave Gunning, Murray Burke

■ *From Shakey the Robot to self-driving cars, from the personal computer to personal assistants on our phones, the Defense Advanced Research Projects Agency (DARPA) has led the development of integrated artificial intelligence (AI) systems for more than half a century. From the earliest days of AI, it was apparent that a robust, generally intelligent system should include a complete set of capabilities: perception, memory, reasoning, learning, planning, and action; and when DARPA initiated AI research in the 1960s, ambitious projects such as Shakey the Robot went after the complete package. As DARPA realized the challenges, they backed away from the ultimate goal of integrated AI and tried to make progress on the individual problems of image understanding, speech and language understanding, knowledge representation and reasoning, planning and decision aids, machine learning, and robotic manipulation. Yet, even as researchers struggled to make progress in these subdisciplines, DARPA periodically resurrected the challenge of integrated intelligent systems and pushed the community to try again. In the 1980s, DARPA's Strategic Computing Initiative took on challenges of integrated AI projects such as the Autonomous Land Vehicle and the Pilot's Associate. These did not succeed, but instead set the stage for the several decades of more siloed research that followed, until it was time to try again. In the 2000s, DARPA took on the integrated AI problem again with its Grand Challenges, which led to the first self-driving cars, and projects such as the Personalized Assistant that Learns, which produced Apple's Siri. These efforts created complex, richly-integrated systems that represented quantum leaps ahead in machine intelligence. The integration of sophisticated capabilities in a fundamental way is the key to general intelligence. This is the story of DARPA's persistent long-term support for this essential premise of AI.*

When one thinks about what it might take to build an intelligent system, it is evident that multiple capabilities will be required. Intelligence is generally considered to be reflected in the ability of a system to learn and understand the world around it, and to deal successfully with new or challenging situations. A closer look at what it might take to accomplish this reveals a surprisingly complex set of abilities that must work together. There are many variations on these themes, but roughly speaking, a robustly intelligent, autonomous agent embedded in the real world will need perceptual capabilities to sense and help interpret external signals and phenomena; a set of beliefs about the world, including itself and other agents, cause and effect, and a host of other things relevant to its survival and success in achieving its goals; a variety of reasoning capabilities to determine implications of its beliefs, understand its environment, plan ahead, solve problems, and so forth; a wide array of learning and adaptation capabilities; the ability to affect the world through action; and, some kind of rich communication mechanism along the lines of natural human language generation and understanding.

  ISSN 0738-4602

These desiderata tend to drive research in different directions, but they are not as separable as listing them makes them appear. The overall intelligence of a system will not lie within a single one of these dimensions alone, but rather in their combination. For example, a huge knowledge base of virtually all facts known to humanity in and of itself would not constitute an intelligent system: Applying that knowledge in real situations and augmenting it to reflect a changing world would be critical for an intelligent system to understand the world around it, be effective in its actions, and respond adaptively to new and unforeseen situations. In other words, intelligence arises from the coordinated aggregation of many different abilities.

While the abilities mentioned dictate nothing specific about how an intelligent system encodes its beliefs or carries out its reasoning — or even that it have separable parts corresponding to the essential capabilities — one thing is clear: To be generally intelligent, it needs to use all of these capabilities in a highly coordinated way. For example, humans learn about concepts and individual entities in a variety of ways, yet we regularly face situations in which very different modes of learning have to coalesce to tackle a new task.

Consider putting together a wooden desk from parts. We would identify the parts through physical interaction with the world: seeing, touching, and so forth. But even perceiving these as separate parts is driven by prior knowledge of physical objects and how physical things can fit together, as well as learned conceptualizations of classes of physical objects such as boards and screws and drawers and legs. Similarly, we'd need to locate and sense the instructions, and find and manipulate tools such as screwdrivers. Putting the object together is not just a matter of perception; background knowledge on the concept of a desk and its uses are needed, for example, in noticing that we are installing a drawer upside down. We may have received verbal instruction on what to do, we may have read about desks and construction in a book, and/or had prior physical experiences working with angles of boards, screws, and delicate pieces of wood. In all cases, we would have had to parse and understand what we were told or read, integrate it into knowledge we previously had about desks and construction projects, and apply it to the current case. This includes generalizing specific knowledge, specializing general rules, and drawing analogies to prior related cases. We'd have to take action through our hands and fingers and make sure that what we intended to happen with an action actually did — and figure out what to do if it didn't. And, we might in the end need to contact another person and ask for help or an assessment of the results of our labor. Even the seemingly simple example of building a desk reveals that none of the abilities of perception, knowledge, reasoning, learning, or communication are alone sufficient to pull off the task, and, in fact, in intelligent systems in the real world, they are deeply intertwined.

Over the years, we have found that the foundational core elements of intelligent systems, from perception to natural language conversation and from reasoning to learning, are each extraordinarily challenging to replicate in machines. Deep and persistent exploration in these lines of research is necessary to begin to create the possibility of human-level performance. Each of these areas of pursuit has developed its own distinctive set of approaches, shared knowledge, and challenging research questions, and over the years the separate research communities have grown apart. We have also learned that even with outstanding progress in each of these directions, the simplistic aggregation of their results in a single system does not yield intelligence. The ways in which the various capabilities make use of one another and depend on one another in a fundamental way is the critical ingredient.

Integrated Systems

Many artificial intelligence (AI) systems have been built to show off or test a particular capability, such as natural language translation or scene understanding. For tasks of real-world complexity, such systems will often include multiple components — for example, a separable knowledge base used for computing the meaning of input sentences or capturing general properties of viewable physical objects. More often than not, the additional components are simply the supporting cast used to bolster the main capability or central algorithm, such as a convolutional neural network algorithm for face recognition. Arguably, most of the successful AI systems that come to mind have been built in this way.

The successes of these systems in their specific realms, while undeniable and impressive, typically exemplify some form of narrow, rather than general, flexible intelligence. We might consider them intelligent, but only in a restricted, limited sense, and we do not take them as examples of the kind of general intelligence that AI as a field seeks as its grand long-term goal. Limited-application systems such as chess-playing programs usually lack a deep, general understanding of even their own domains — they don't even know they are playing chess, that chess is a game, and that their moves are expressly in response to an adversary whose goal it is to defeat them — let alone a broader view of the world beyond their own constricted lanes. They do not have a capability that we generally think of as central to a generally intelligent system: the ability to be coached to take on different sets of goals or tackle significantly different kinds of tasks. And such systems are typically not resilient when the world presents unanticipated situations. Narrow AI systems tend to be fragile, sometimes with surprisingly hard boundaries on their capabilities, leading to catastrophic, as opposed to graceful, failures when they are stumped. No matter how much a narrow, specialized AI system excels in its own

DARPA Changed My Life and the Trajectory of Google

Have you heard of self-driving cars? Of course you have! This nascent technology carries the promise of saving one million lives every year. It will let blind people and children get around effortlessly. And for many of us, it'll turn our daily commute into pleasure time for watching movies or taking a nap. Or — God forbid — indulging in WhatsApp and Instagram.

This all would not be without the Defense Advanced Research Projects Agency (DARPA). In 2003, DARPA launched a prize competition known as the Grand Challenge: Who could build a car that navigates for over 130 miles of desert trails without a human driver?

In 2004, no robot survived for 30 minutes.

A year later, five robots made it across the finishing line. The route involved steep inclines and treacherous mountain switchbacks. At Stanford University, we had entered Stanley into the race. Thanks to our emphasis on machine learning and adaptive intelligence, Stanley won. And fortunately for the world, Google cofounder Larry Page was in the crowd.

A few years later, Page convinced me to start the Google Self-Driving Car Project, now known as Waymo. Most of my early team members were veterans from the Grand Challenge. Highlighting just three: Chris Urmson, who directed the Carnegie Mellon University team that later won the DARPA Urban Challenge; Michael Montemerlo, Stanley's indefatigable software chief; and, Anthony Levandowski, who built the world's first self-riding motorcycle.

It is hard to overstate DARPA's contributions. While Stanley was a far cry from the amazing work done by the Waymo team today, self-driving cars would not exist without the vision and leadership of DARPA, specifically Director Tony Tether and Program Manager Ron Kurjanowicz. The DARPA Grand Challenge was the birthplace of the modern self-driving car. It earned Stanley a place in the Smithsonian Air and Space Museum.

And I hope it will change your life, and the trajectories of generations to come.

– Sebastian Thrun

lane of expertise, it would fall short of our intuitive view that an intelligent system should understand its world and adapt successfully to new or trying situations.

To achieve more generally intelligent systems, we need the individual realms of competence within a system to be integrated. It makes sense to build the very best vision component we can, and live with the fact that its technology is not likely to be equally capable of complex planning. But in a robust, generally intelligent system, that vision piece cannot keep fully to itself. Consider the visual input channel for a mobile robot: The robot's success relies upon a very tight feedback loop between its perceptual apparatus and its general background knowledge, such that objects, events, and actions in the external world can be interpreted in a way that can feed the robot's decision and planning components. Similarly, when the robot takes an action, it must perceive (perhaps through multiple perceptual channels) the results of its action to plan its next steps accordingly. While the algorithm for planning an action or inferring the consequences of its effects is likely to be very different than one for processing a low-level visual signal, an effective perceptual subsystem of an intelligent system must be intimately connected to, and able to communicate with, the other subsystems.

Similarly, intelligent machines typically will not have a separate learning component. In the natural world, learning is richly — perhaps even indistinguishably — interconnected with perception and reasoning. Humans regularly remember things they've seen and done, and integrate those experiences with many different types of knowledge they have stored from prior experiences. While we can preload a huge amount of knowledge into a machine, it may not be the same as knowledge a human has acquired through learning. Knowledge is also dynamic, being regularly augmented and corrected based on the ongoing stream of experiences. Even though machine learning, representation, and reasoning have generally grown up as separate subfields of AI, and are each given their own article in this collection, a robustly intelligent system embedded in the real world will have these capabilities entirely intertwined.

These capabilities imply that the overall architecture of a generally intelligent system is paramount. A block diagram may show different boxes for syntactic processing, visual object recognition, generalization inferences, and planning, but the sharing of their processing and the ways they influence the behavior of each other is as important as the boxes themselves. In fact, this is a challenge of using standard architecture diagrams to attempt to capture system structure: They do not elucidate the rich, complex ways in which the components need to share information, work closely together, and influence one another.

Many AI scientists have looked to cognitive science and psychology for inspiration on possible interconnected architectures; the notion of a cognitive architecture has hidden behind many AI system implementations, even if not called that. From the earliest days of groundbreaking AI scientists such as Allen Newell, optimal ways of piecing together the parts of an intelligent system have been the subject of much research. Newell and other cognitive psychologists have developed a variety of computational models that strive to capture the major functions of human cognition, such as Adaptive Control of Thought—Rational and Soar (Laird, Lebiere, and Rosenbloom 2017). These typically represent human intelligent processing in terms of perception, cognition, memory and learning, and action selection, often with some account for attention.

The space of cognitive architectures explored by the research community is immense and rich; Kotsuruba and Tsotsos (2018) catalog 84 architectures culled from an investigation of more than 900 projects. The variations are interesting and provocative, but for our purposes the key is the way in which cognitive architectures attempt to integrate the various components that might make an intelligent system. Consider a very simple scheme for building an intelligent system that we used in DARPA's Information Processing Technology Office (IPTO) as a framework for focusing integrated intelligent system projects

Figure 1 illustrates the view that we used internally in IPTO as we were developing our Cognitive Information Processing Technology Initiative launched in 2002. The main ideas are simple. Active perceptual mechanisms are well-connected to background knowledge and imbued with the ability to adapt and learn (through deliberative and reflective processes); they are the interface between the brains of the AI system and the signals from the outside world. This implies that prior knowledge and even a certain amount of reasoning may influence how objects, scenes, and actions may be perceived and understood. Language input from other agents in the world (received through the perceptual layer) illustrates how these capabilities may need to be very sophisticated and tightly integrated with other types of processing. Effectors of various sorts, such as wheels for vehicles, hands for humanoid robots, speech synthesizers, and so forth, are the mechanisms for the intelligent system to take action in the world. Spoken language generated by the AI system gets out into the world through this pathway. Internally, it seems plausible to posit artificial system operation at multiple levels of deliberation, as in humans and some animals: some interpretation and subsequent action is automatic and purely reactive; a great deal of internal processing is deliberative in that reasoning is involved (for example mental simulation and planning, case reasoning, and diagnosis); and some processing is what we might call reflective, in that it takes internal observation of the agent's own reasoning and behavior into account.

This is clearly a simplified view and not necessarily a prescription for how to build an intelligent system. Note, in particular, that it does not posit a single, separate learning component, but instead implies that learning mechanisms will be distributed throughout the architecture. It provides a useful framework for discussing what elements are critical to the implementation of a generally intelligent system and offers a way to visualize which capabilities need to share information, and to a degree how they share it. In much of the integrated intelligent systems work supported by DARPA (historical programs are highlighted in figure 2), program managers looked for consolidated architectural views to ascertain that integrated AI projects came not only with a view of the core components, but also with an understanding of the critical relationships between those components and how they shared data, knowledge, and processing.

Finally, it should be stated that standard approaches to systems integration, meant to bring together multiple pieces of software and hardware, do not attack the integration problem at the correct level of abstraction for building an intelligent system. While such approaches may indeed be important for connecting the software and/or hardware pieces of an implemented system, the kind of integration we focus on here is concerned with information, knowledge, and skills, intimately connected with learning and adaptation. Multiple learning algorithms within the same system must be coordinated to ensure that their impact on stored knowledge (that is, long-term memory) is aligned, compatible, and consistent. Beliefs and intentions from multiple sources across the system must be integrated and reconciled to be brought to bear on a situation that needs a response. This all has more to do with the alignment of knowledge/beliefs than with interfaces between components.

Two Key Threads, Under DARPA's Leadership

Given the natural tendency for technical subcommunities to focus on their specialized lines of inquiry, extra incentive is usually needed for scientists to focus on serious integration within a multicomponent intelligent system. The impetus usually has to come from an external source, such as a project sponsor or funding agency, through a significant large and broad research challenge. It takes a special challenge — and sizable funding — to tackle a broad-based system seeking to embody more general intelligence.

DARPA is unique in its ability to challenge broad communities with ambitious endeavors, the success of which require fundamental science and engineering

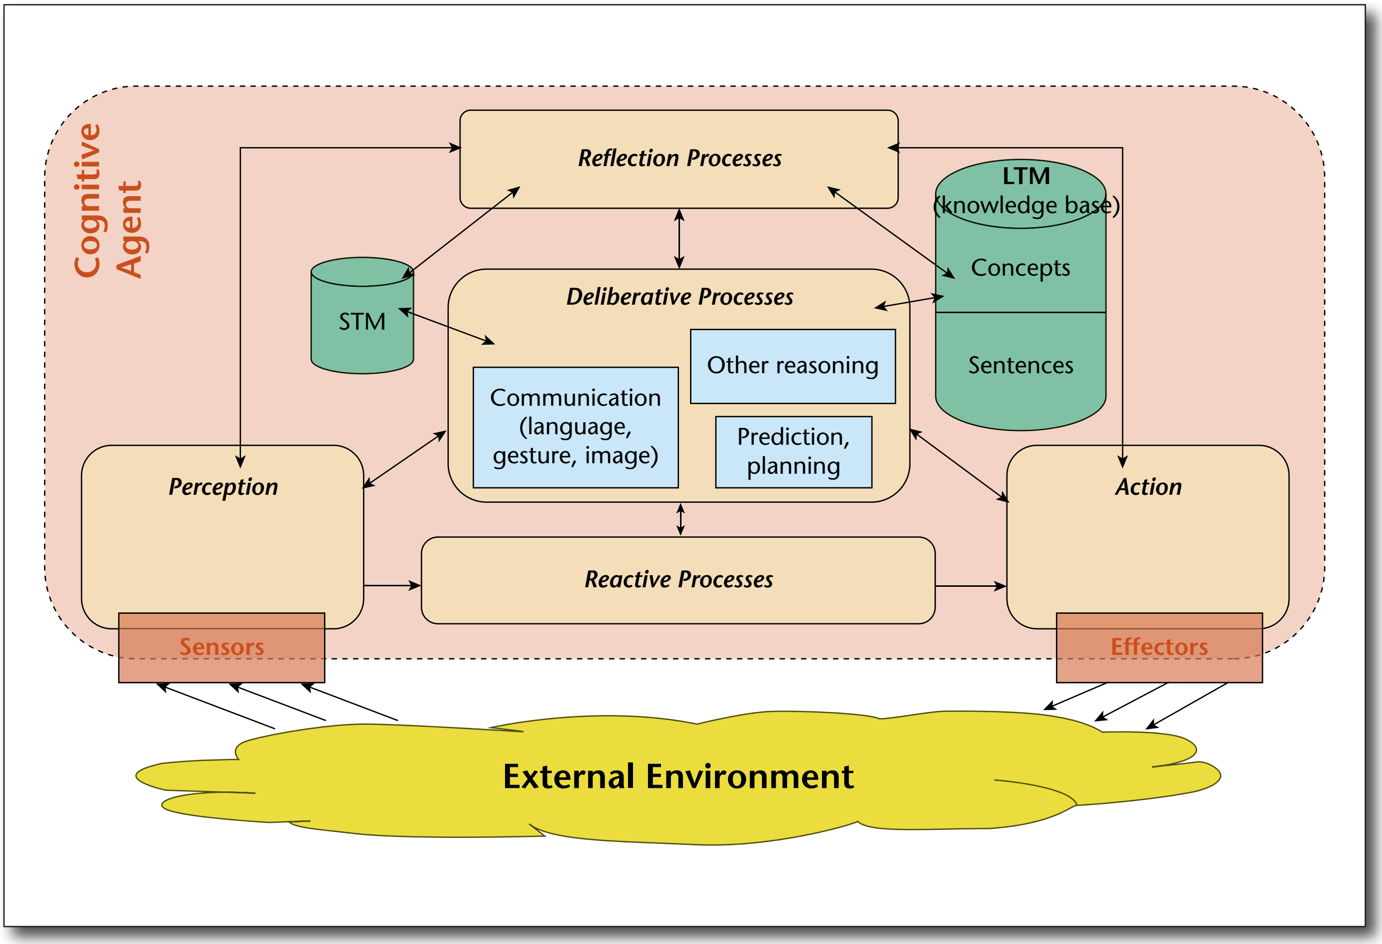

Figure 1. Example of Cognitive Architecture.

Figure courtesy of DARPA.

breakthroughs. AI has always been one of DARPA's preeminent interests, but the Agency has distinguished itself in driving the research community to tackle full-scale integrated AI systems. By funding sizable ambitious efforts — often with large, multiorganizational teams — DARPA has been the key player in spurring the community to develop integrated intelligent systems.

DARPA's persistent, long-term efforts to guide the community toward integrated AI systems has largely, and consistently, been along two paths. While more active in some time periods than others, the Agency has focused the energy of the AI community toward autonomous mobile platforms and personal assistants (or partners) for intellectual tasks. In the earliest days, the mobile platform focus was on self-guided robots that might perceive their environment, plan their next moves, and react appropriately to environmental changes, while achieving very simple goals, such as moving boxes and going through doorways. This thread eventually led to the development of semiautonomous and fully autonomous vehicles of which there is now broad public awareness. On the intellectual partners front, things started with a dream of human–machine symbiosis and memory-aiding assistants, and evolved over 45 years to personal assistants on mobile phones and home-based devices powered by Apple's Siri, Amazon's Alexa, and others. These technologies are now taken for granted and, in hindsight, it seems as though it was always inevitable that smart, autonomous vehicles and intelligent personal assistants would have places in early 21st century society. But, they may never have come into being without the extraordinary foresight and ambition provided by DARPA's leaders from as far back as the 1960s, and its persistent, dedicated pursuit over the intervening decades.

Autonomous Mobile Platforms

Robots and autonomous vehicles have been a subject of fascination since at least the time of Leonardo Da Vinci, who conceived of a robot in German knight's armor in 1495. His idea was rediscovered in the 1950s when his enormously complex notes were deciphered. However, there is no evidence that Da Vinci actually constructed his robot. More modern excitement about humanoid robots was presaged by Electro, displayed by Westinghouse in the 1939

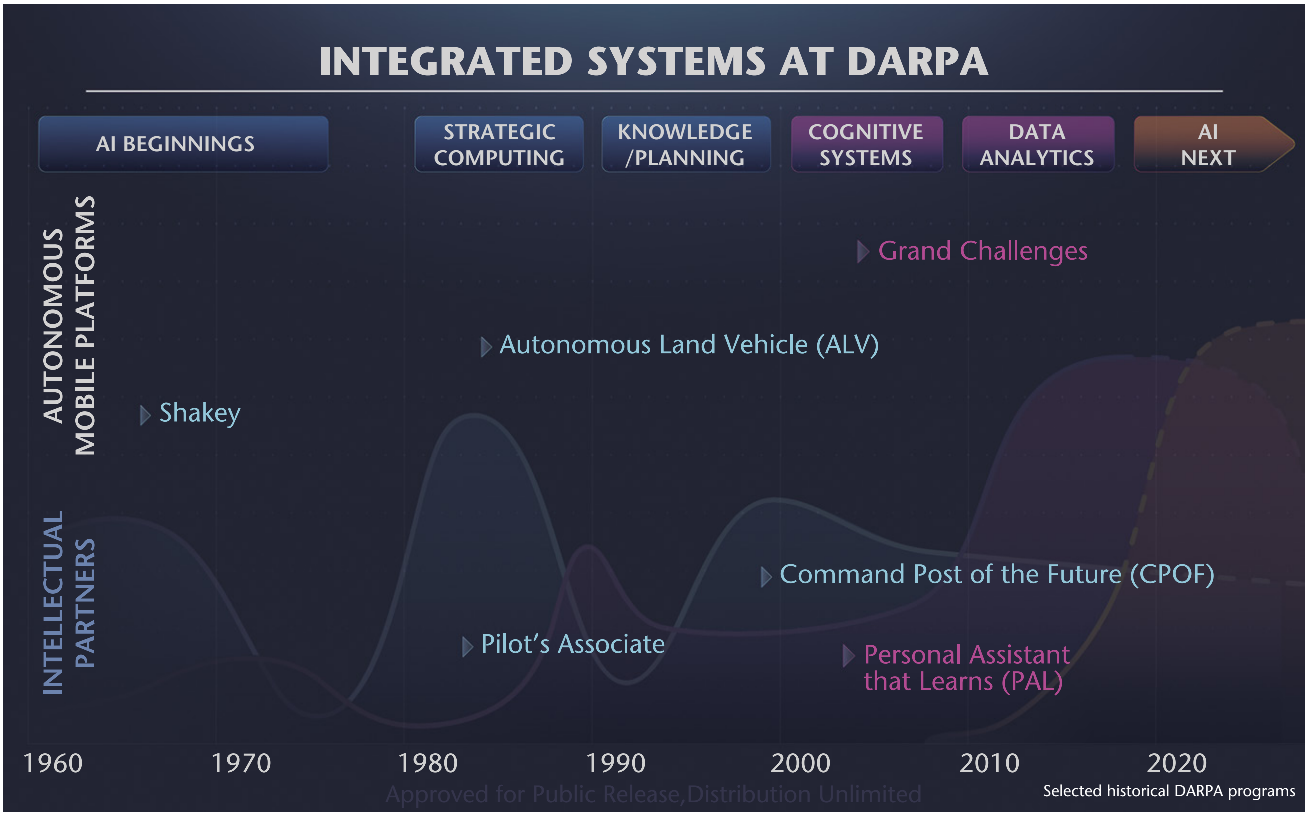

Figure 2. Integrated Systems at DARPA.

The historical DARPA programs highlighted in this article. *Figure courtesy of DARPA*.

New York World's Fair; people would wait in line for three hours to see the 7-foot–tall, 250-pound Elec- tro. It walked (rolled), talked, and smoked. It was commanded by a telephone handset with a reported vocabulary of over 700 words that played via a 78-RPM record player near its mouth. Beyond crude early attempts to build robots, of course, modern science fiction is replete with portrayals of robots, androids, robot/human hybrids, and vehicles that can think for themselves.

DARPA has had a continuous interest in autonomous robot and vehicle research and development (R&D). Three representative seminal programs are Shakey the Robot in the 1960s, the Strategic Computing Initiative's Autonomous Land Vehicle (ALV) program in the 1980s, and the DARPA Grand Challenges in the 2000s. Each reflected major advancements in component technologies at the time and demanded new and significant developments in autonomous systems capabilities via integration. Indeed, each served as a driver for the research and development community to attack hard technical problems and tackle the kind of integration of complex technologies that is necessary for autonomy to become a reality.

The complexity of the real-world challenges in these robotic contexts made for a compelling and complex integration need. Control of multiple sensors was required to detect distances, terrain, and impediments. Sensor outputs needed to be combined, which, in turn, necessitated reasoning about inconsistent input; potential obstacles to avoid, move, or otherwise work around; actuators to sequence and control; and situation models to be updated as the platform moved. There were multiple continuous observe, decide, and act cycles requiring coordination until the vehicle reached its goal (destination).

Shakey the Robot

In 1966, one of DARPA's long-term AI R&D goals was the development of automatons capable of gathering, processing, and transmitting battlefield information. A shorter-term goal was to develop a mobile, autonomous agent that could accomplish nontrivial tasks in a real environment. From 1966 to 1972, SRI

International (then called the Stanford Research Institute) conducted research on the first general-purpose intelligent mobile robot, nicknamed Shakey, because it shook when it was stopped. Shakey, shown in figure 3, was the first robot able to perceive and reason about its own surroundings, as well as reason about its own actions.

Shakey's limited world contained a number of rooms connected by corridors, along with doors and light switches with which the robot could interact. Shakey had a short list of available actions within its repertoire including traveling from one location to another, turning light switches on and off, opening and closing doors, rolling up and down ramps, and moving objects.

Shakey's software integration was critical for it to achieve its goals and can be viewed as a three-layer hierarchy. The lowest level consisted of the control programs that drove all the motors, controlled platform movements such as ROLL and TILT, and captured sensory information. The middle level supervised the primitive actions it could carry out and comprised prepackaged combinations of the control programs to accomplish higher-level actions such as GOTO or PUSH. This level also included the vision processing algorithms associated with the camera. Solutions to problems were planned within the top layer, ultimately using the Stanford Research Institute Problem Solver (STRIPS) planning and inference mechanism. This was the layer in which SRI implemented and integrated new search heuristics, action control, and planning functions, as well as learning and scene understanding. Overall, this was a simple implementation of the kind of cognitive architecture we introduced earlier.

Shakey's world was represented in a grid model and probably was the first to use adaptive cell decomposition for robot motion planning. The problem of locating and navigating the shortest path from source to destination resulted in SRI's invention and integration of the A* algorithm, a milestone in AI history (Hart, Nilsson, and Raphael 1968). STRIPS would deduce a sequence of goal-satisfying actions from initial conditions in a backward-chaining fashion. A typical STRIPS plan for Shakey was six steps. Plan learning required storing the previously deduced series of steps and allowing a macro-operator to be used whenever preconditions were appropriate. Shakey did not attempt complete visual scene analysis but rather concentrated on acquiring specific information needed to perform its tasks. This required a tight integration between the planning and vision systems. Around 1970, SRI integrated a parser and semantic analyzer with STRIPS that translated commands in simple English to STRIPS logical statements. SRI also developed and integrated axioms to enable Shakey to deduce solutions to more indirect problems such as pushing a box off of an elevated platform. The robot needed to infer that it had to roll up a ramp first and push an intervening obstacle out of the way before accomplishing the goal. This necessitated an integrated control loop including its movement programs, STRIPS inference software, and its bumper sensor control software.

Funding for Shakey wrapped-up in 1972. DARPA remained interested in the AI reasoning challenges but not the robotic hardware problems. More detail on other DARPA robotic R&D is provided in the Vision and Robotics companion article (XXXXX 2020). Shakey is currently on display at the Computer History Museum in Mountain View, California. In 2017, Shakey was designated as an "IEEE Milestone in Electrical Engineering and Computing" (see Engineering and Technology History WIKI 2017) for IEEE's description of the historical significance of the work).

ALV Program

After a quiet dozen years on the mobile platform front, a major new impetus came from DARPA's Strategic Computing Initiative, in the form of the ALV, a major applications program that began in August of 1984 (see figure 3 for a photo). The US Army Engineering Topographic Laboratory helped to oversee the program and its goals were aligned with the Army's long-range strategic vision of using autonomous vehicles in logistics and supply, search and rescue, and possibly combat operations.

The Martin Marietta company was selected as the integration contractor with an initial funding of $10.6 million for a period of 42 months. Subcontractors for R&D and subsystem elements included SRI, Carnegie Mellon University, the University of Maryland, Hughes Research Laboratories, Advanced Decision Systems, and the Environmental Research Institute of Michigan. ALV's performance objectives related to driving within an unconstrained environment, which required tight integration of the following subsystems developed by the subcontractors: Sensors (the ALV sensor subsystem used an RCA color-video charge-coupled device TV camera and an Environmental Research Institute of Michigan-designed laser range scanner), Perception (the perception subsystem accepted sensor images to define road edges and produced coherent scene models), Reasoning (the reasoning subsystem received a plan script from a human test conductor and coordinated all ALV actions; it obtained scene models from the perception subsystem and converted them into smooth trajectories that were passed to the pilot subsystem to drive the vehicle), Pilot (the pilot subsystem converted the intervals of a trajectory into steering commands for the vehicle; it calculated steer-right, steer-left, and speed commands), Knowledge Base (the knowledge base consisted of a digital representation of the road network and nearby terrain), Vehicle (the vehicle subsystem had an undercarriage that was an eight-wheel hydrostatically driven unit capable of traversing rough terrain at speeds up to 29 km/hour and 72 km/hour on improved surfaces; steering was accomplished by reducing or reversing power to one of the wheel sets), and Human Interface (the human test conductor directly entered the plan script for the road-following test; a dead-man

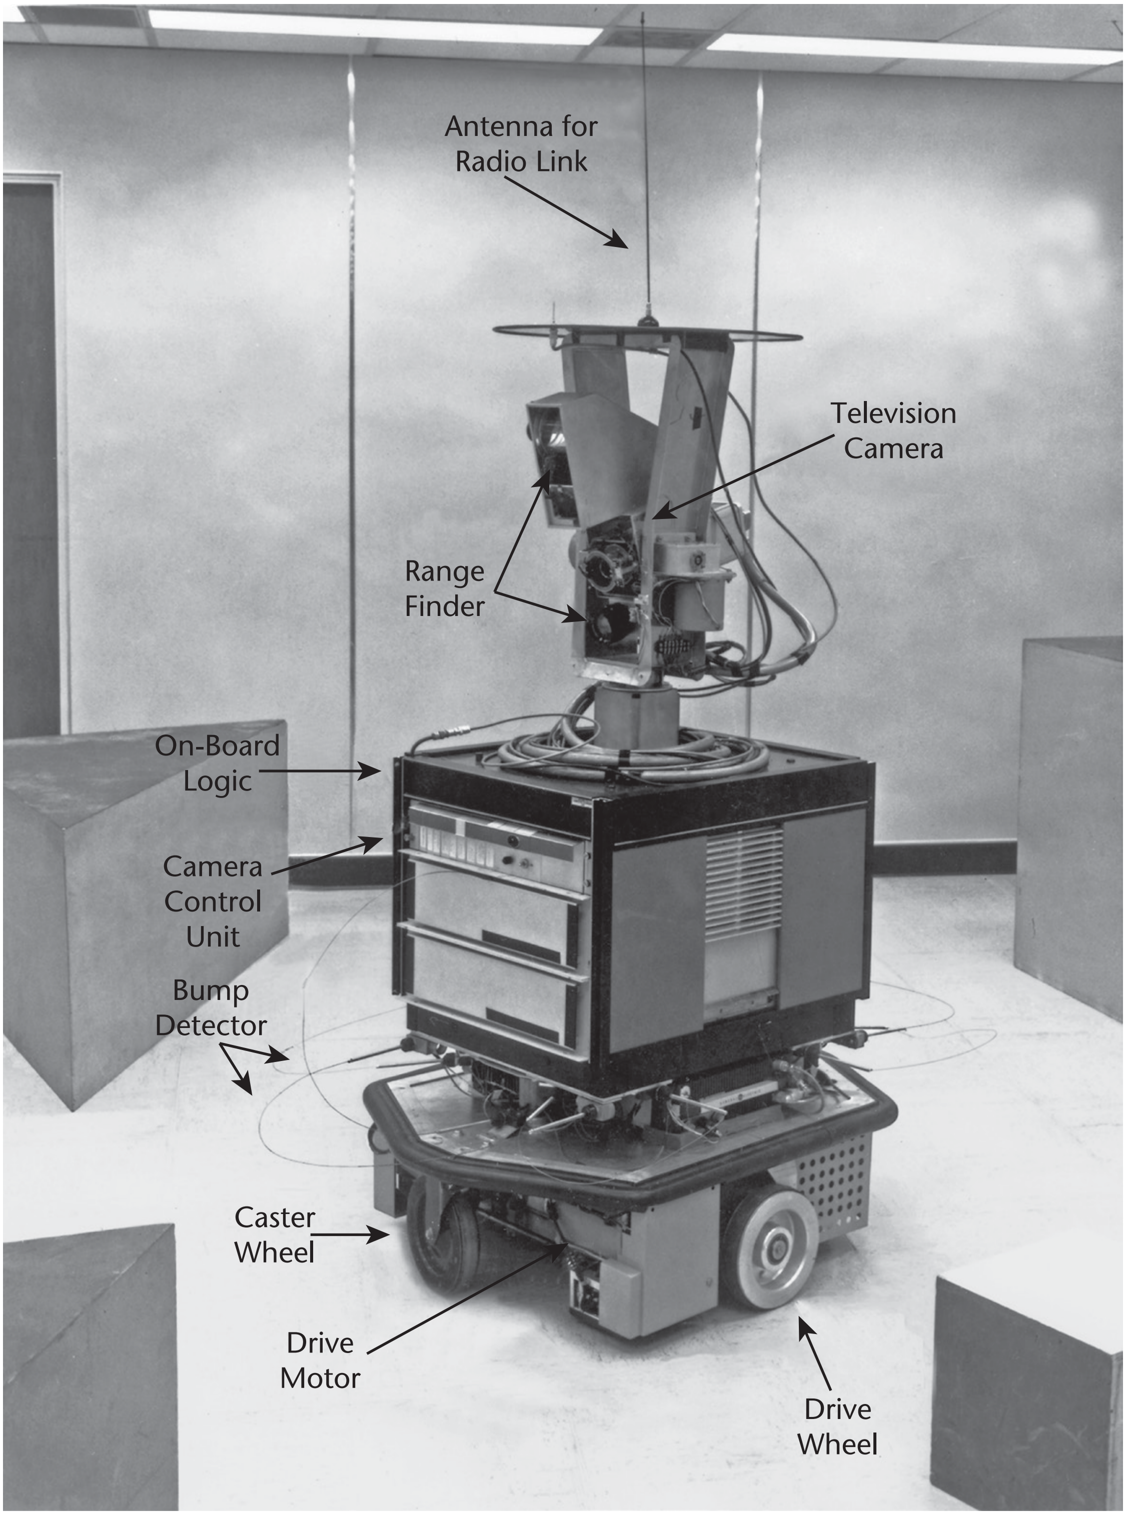

Figure 3. Shakey the Robot.

Figure adapted and reproduced with permission from SRI International.

switch served as a safety device for halting unexpected or out-of-control trajectories).

In August 1987, the ALV performed the first autonomous cross-country traverse based on sensor data. During this and subsequent trials extending for about a year, the ALV navigated around various kinds of isolated positive obstacles over traverses of several kilometers. The vehicle reached speeds of 5 km/hour and completed about 75 percent of the test traverses. Processing the image data accurately and quickly was the major challenge in these tests. The ALV is depicted in figure 4.

DARPA cancelled funding for the ALV in 1988, but interest in autonomous vehicles continued at DARPA as well as in academia and industrial R&D laboratories.

DARPA Grand Challenges

After a decade and a half of research on the individual technologies (for example, sensors, perception, machine learning, robotic control, reasoning, and planning), DARPA took another ambitious run at the problem of autonomous driving with the DARPA Grand Challenges in 2004 and 2005, and the Urban Challenge in 2007. Before the Challenges began in 2004, self-driving cars were a science-fiction fantasy. By the end of the Urban Challenge in 2007, DARPA had demonstrated the feasibility of autonomous driving, even in traffic, and set the stage for industry investment to turn this high-risk technology into reality.

The first DARPA Grand Challenge attracted 106 applications from teams wishing to compete for a $1 million prize. The goal was for a vehicle, unaided by humans, to drive itself over a more than 140-mile course of complex, natural terrain in the Mojave Desert from Barstow, California to Primm, Nevada. In the end, none of the 15 finalists made it farther than nine miles into the 142-mile desert course. At that point, the challenge seemed impossible and to some, DARPA appeared foolish for taking on such an ambitious goal. However, 18 months later, in the second DARPA Grand Challenge, four teams of an original 195 completed a 132-mile desert course in Nevada with a Stanford University team taking what had become a $2 million prize for the winning performance of its vehicle, Stanley. In 2007, just two years later, six teams completed the Urban Challenge, a second follow-on competition to add vehicle traffic in an urban environment to the challenge. With its vehicle, Boss, a team from Carnegie Mellon University in Pittsburgh crossed the finish line first and won the $2 million prize. The second and third place teams from Stanford and Virginia Tech took home $1 million and $500,000, respectively.

The real turning point in this story came in the second Grand Challenge in 2005, when six teams successfully finished the challenge and this impossible task suddenly became doable. Taking a more detailed look at Stanley, Stanford's entry in the challenge that finished in first place, Sebastian Thrun, the leader of Stanford's Stanley team, recounts the team's approach to the Grand Challenge (Thrun et al. 2006):

> In relation to previous work on robotics architectures, Stanley's software architecture (figure 5) is related to the well-known three-layer architecture (Gat et al. 1998), albeit without a long-term symbolic planning method. A number of guiding principles proved essential in the design of the software architecture: from a broad perspective, Stanley's software mirrors common methodology in autonomous vehicle control. However, many of the individual modules relied on state-of-the-art AI techniques. The pervasive use of machine learning, both ahead and during the race, made Stanley robust and precise. We believe that those techniques, along with the extensive testing that took place, contributed significantly to Stanley's success in this race.

Because of DARPA's embrace of autonomous robots and vehicles, the world has come to accept autonomous vehicles as a realizable vision and, as a result, we will see radical societal and military changes ramp-up in the very near future. Shakey, the ALV, and the DARPA Grand Challenges illustrate the progression of autonomous systems capabilities from primitive prototypes and demonstrations to an emerging military and commercial application market. As the DARPA programs demonstrated, the functioning of any autonomous system in the real world necessitates the tight integration of perceptual, cognitive, and action mechanisms. The integration tasks are complex and costly, but autonomy emerges when integration is successfully accomplished.

Intellectual Partners

From the invention of personal computing to the creation of the first personal assistant on your phone, DARPA has led the development of integrated AI systems capable of acting as intelligent partners for humans. In the 1960s, DARPA formulated big visions and initial prototypes for personal computing and personal assistants. In the 1980s, following AI's first and worst AI funding "winter," DARPA took on an ambitious effort to develop the Pilot's Associate, a personal assistant for fighter aircraft. In the 1990s, DARPA's goals focused on personal Anchor Desks and decision aids for commanders and their staff, such as the Command Post of the Future (CPOF). In the 2000s, DARPA again went all out with an ambitious thrust in cognitive systems that included the Personalized Assistant that Learns (PAL) program, which led to Apple's Siri. Today, DARPA is continuing to push beyond the state-of-the-art with programs such as Communicating with Computers, which has the goal of enabling computers to act as full communication partners with human beings.

Visionary Concepts

In 1963, J.C.R. Licklider became the first Director of DARPA's IPTO. Licklider built upon Vannevar Bush's concept for a MEMEX, a personal assistant for intellectual rather than physical activity (Bush 1945), to formulate his vision for personal computing and Man–Machine Symbiosis (Licklider 1960). IPTO began funding a number of projects to pursue that vision, including sponsoring Doug Engelbart to formulate his ideas for Augmenting Human Intellect (Engelbart 1962). The

Figure 4. Autonomous Land Vehicle.

Figure courtesy of the Lockheed Martin Corporation.

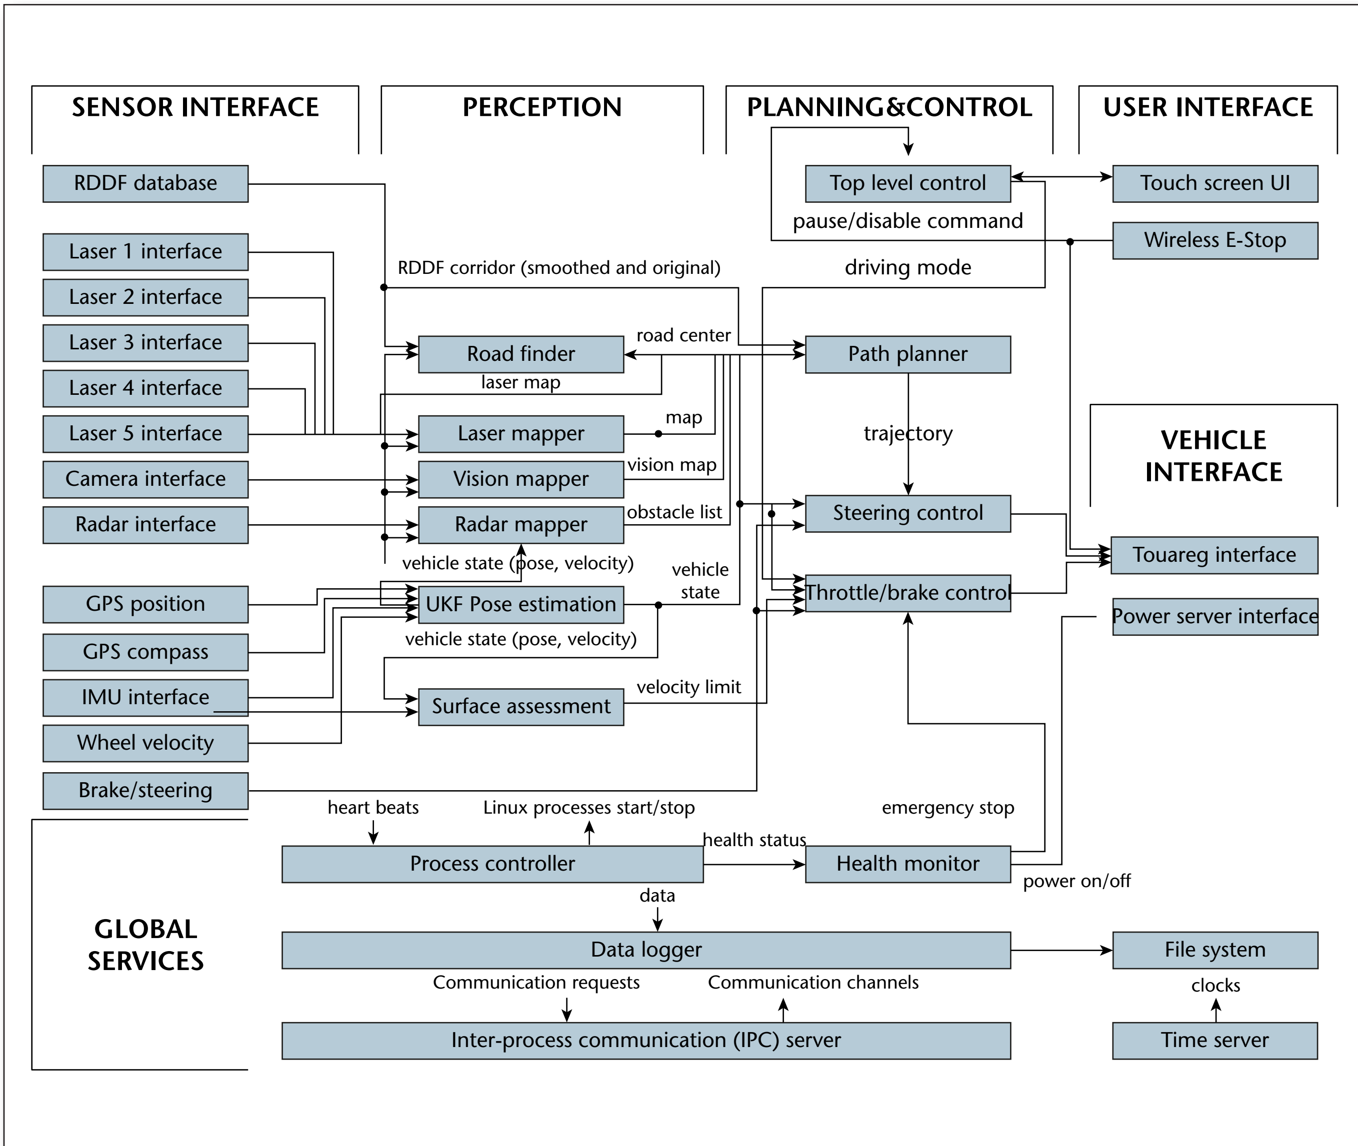

Figure 5. Flowchart of the Stanley Software System.

The software is roughly divided into six main functional groups: Sensor Interface, Perception, Planning & Control, User Interface, Vehicle Interface, and Global Services (being a number of cross-cutting services, such as the process controller and the logging modules). *Reproduced with permission from Thrun et al.* (*2006*). *Copyright (2006) Wiley Periodicals, Inc., A Wiley Company.*

goal of the effort was to develop an intellectual assistant that would augment human ability to understand and solve complex problems. In 1968, Engelbart gave the first demonstration of this early concept with what came to be called the "Mother of all Demos," which included an amazing integration of the first mouse, hypertext, word processing, and distributed collaboration. This integrated system was designed to complement human cognition and collaboration with the mouse as a new way for a machine to perceive manual input; hypertext to mirror human semantic memory; and a set of actions to edit and share semantic knowledge.

From this initial concept demonstration, DARPA pursued a series of programs that further developed the vision for an intelligent personal assistant. First was the Computer-Based Consultant in the 1970s (Nilsson 1974), which hoped to provide advice to maintenance technicians. This included initial elements of an integrated system: a vision component, a Hierarchical Planner (Sacerdotti's Noah [Sacerdoti 1980]), and a Natural Language Processing module (Task Dialog Understanding System from Barbara Grosz [Grosz 1974]), which included a grammar, called DIAGRAM, for representing and reasoning about processes and goals. The project was canceled before completion, at the onset of AI's first funding winter. But it was not the end of integrated AI systems.

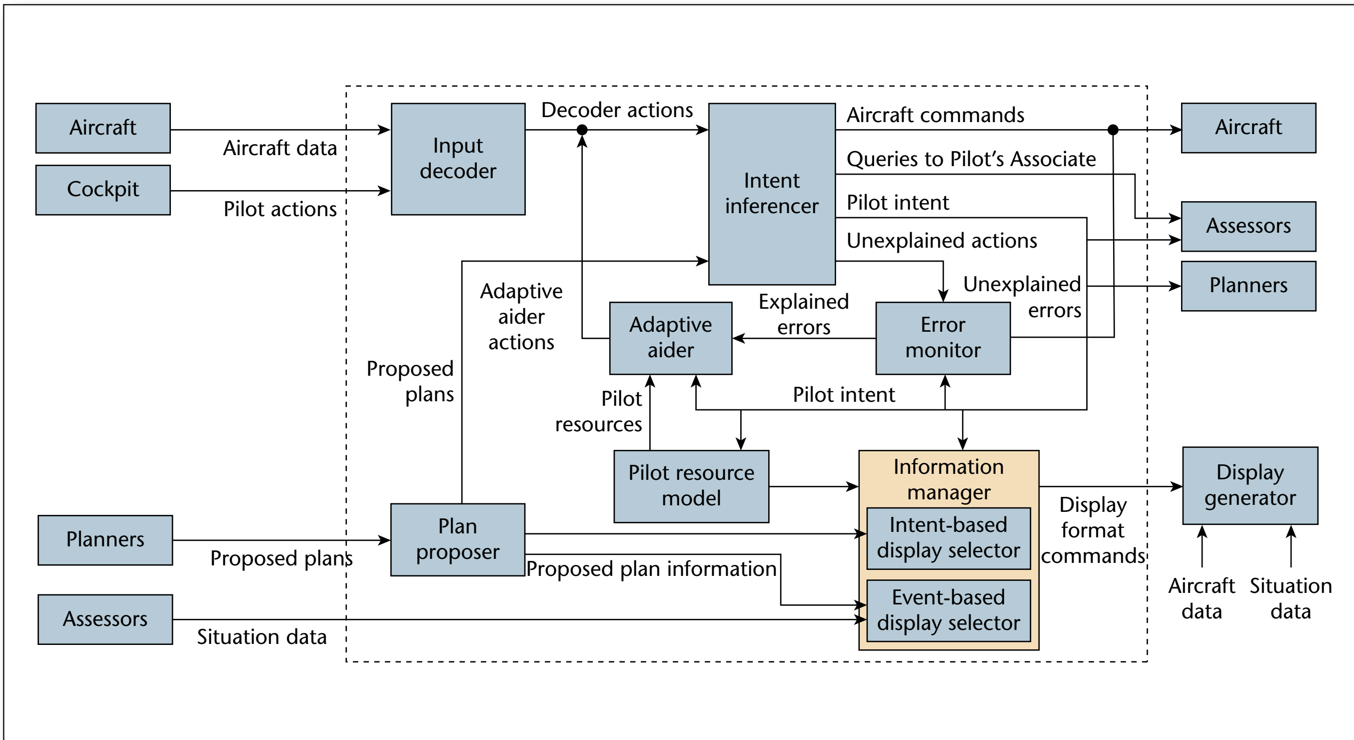

Figure 6. Flowchart of the Demo 3 Pilot–Vehicle Interface.

©1991 IEEE. Reprinted, with permission, from Banks and Lizza (1991).

Pilot's Associate Program

As the funding winter thawed in the next decade, DARPA initiated one of its most ambitious efforts, the Pilot's Associate (Banks and Lizza 1991). The goal of the Pilot's Associate was an integrated assistant for fighter aircraft, essentially a virtual backseater. It would take in sensor data from the aircraft, maintain a model of the situation, develop plans for tactical maneuvers and weapons operations, and assist the pilot in task execution. The integrated architecture (see figure 6) included cooperating knowledge-based subsystems for situation assessment and system status, with a tactics planner, mission planner, and pilot–vehicle interface. The pilot–vehicle interface would infer the pilot's intent, which would allow the right information to be presented the pilot and appropriate supporting actions to be taken or suggested. The program had a number of impressive accomplishments, and also ran into a number of unanticipated issues that prevented delivery of the capability.

One of the accomplishments involved new insights into integrated AI systems engineering. The Lockheed Pilot's Associate team encountered a very interesting issue at the initial integration of its five expert systems. Even though the team had used detailed system engineering techniques to ensure that each of the expert systems was communicating effectively, when the team finally brought them all together, there were serious issues in terms of the resulting system behaviors. Essentially each expert system was pursuing its own strategy and the result was that they were providing conflicting advice to the pilot. To address this challenge the team came up with the notion of a plan and goal graph, which was a mechanism for sharing contextual awareness across the human/machine system. The plan and goal graph provided a context that coordinated machine planning with pilot intent and pilot actions, and unified the focus of all the system components on the information requirements and action decisions necessary to support a common set of goals and strategy. An example of a component technology shortfall that prevented this capability from being delivered to the warfighters was the lack of a general approach to geometric reasoning (for example, recognizing a maneuver and projecting the geometry that would result).

The impact of the Pilot's Associate program is best summarized by the concluding remarks of the program managers in their 1991 summary report:

> The actual implementation of a system of this complexity uncovers many gaps in technology still to be addressed by the research community. (Banks and Lizza 1991)

CPOF Program

The Pilot's Associate was an extremely ambitious effort, and the lack of a big success could be argued to have started a milder AI Winter. Research funding did not

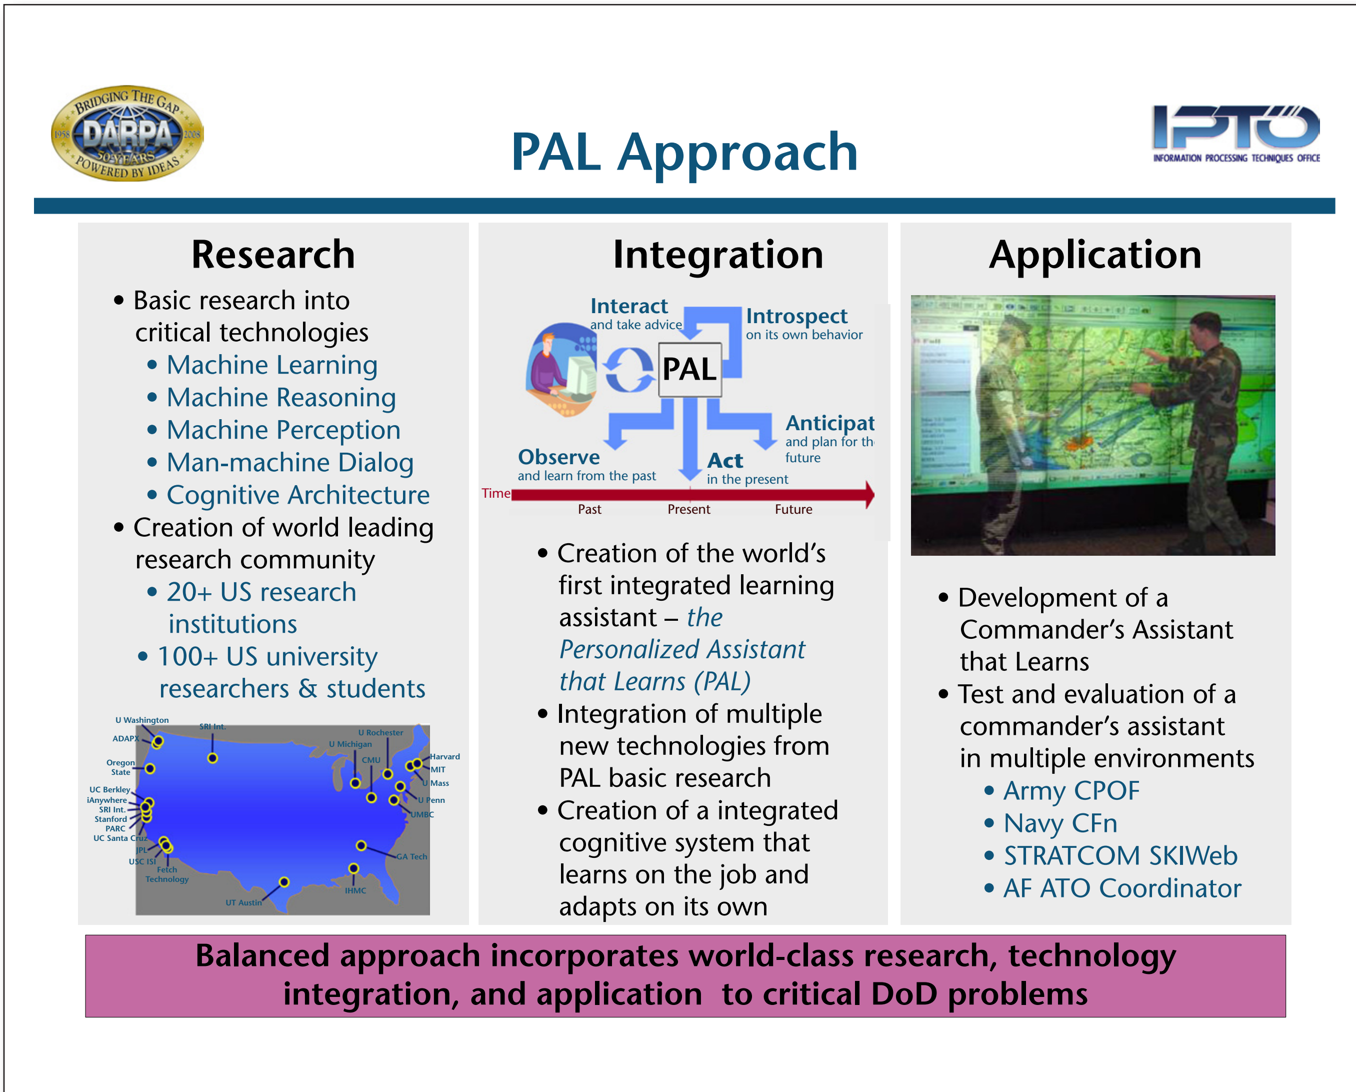

Figure 7. PAL Approach.

Figure courtesy of DARPA.

disappear, but DARPA pulled back from grand AI systems visions and began to focus on individual technologies. Nevertheless, during this time work continued on planning and decision aids, primarily for military command and control: anchor desks, logistics planning tools, and command-and-control environments. Under the umbrella of the ARPA-Rome Labs Planning Initiative, a number of programs pursued the development of planning and decision assistants for commanders and their staff.

This line of work culminated in the CPOF program in the late 1990s, which was eventually transitioned to the Army as their command-and-control system. CPOF began as an ambitious DARPA research effort to build a wide range of components and combine them into an integrated planning and execution environment for commanders and their staff. The initial architecture consisted of modules for information integration, knowledge-base management, planning and reasoning, speech and gesture interaction, and tailored visualization. A key strategy for the CPOF development was a double-helix development strategy (what today we would call agile development) — continuous iterative prototyping and joint development between visionary military users and technologists who would jointly evolve both the technological design and the military concept of operations. The resulting system pruned the portfolio of initial technologies down to a tightly integrated collaborative visualization environment. CPOF in many ways realized Doug Englebart's vision for augmented cognition and distributed collaboration, in the command center.

As is often the case, the final transition included only a small subset of the initially developed technology.

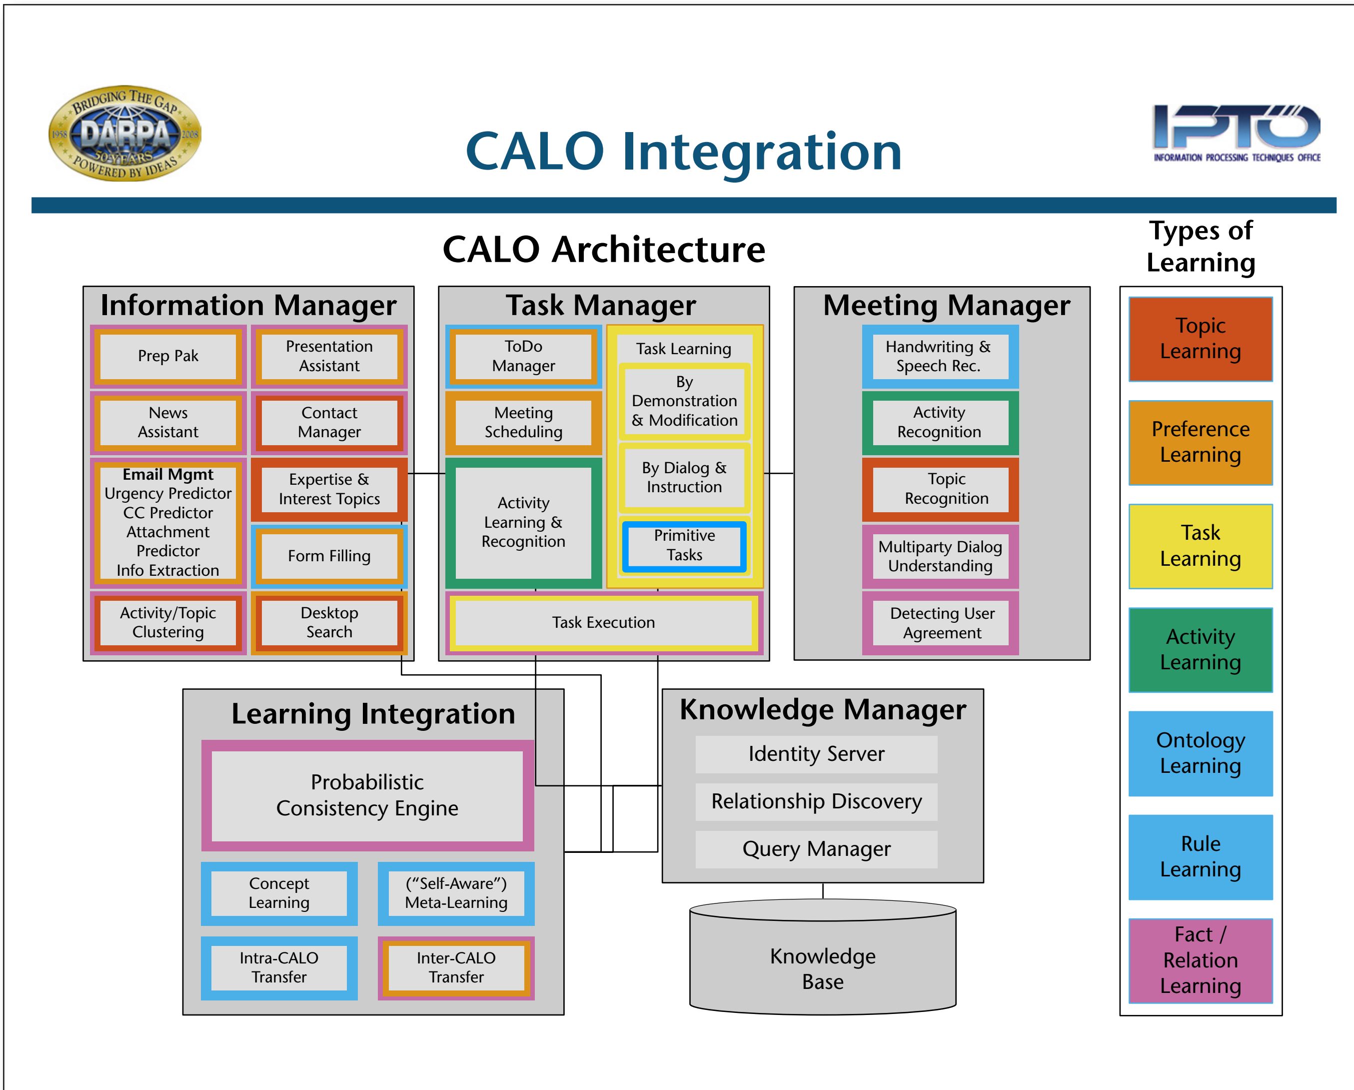

Figure 8. CALO Integration.

Figure courtesy of DARPA.

The heavier-weight knowledge-based components did not make the cut, but lighter-weight, semistructured, schema-less data structures proved flexible and useful for applications. The system that was transitioned was able to access and integrate information from dozens of stovepiped information sources to create dynamic integrated, but individualized, views of the battlespace to support collaboration. Individual users could drag and drop important elements from each other's workspace to create an extremely powerful and agile collaboration capability. This was an outcome of the double-helix prototyping process and was enabled by a lightweight schema-less data structure integration infrastructure (Chapman et al. 2005). The full CPOF system included speech and gesture interaction, but this technology was also dropped in the double-helix process and had to wait until the PAL program, several years later, to be ready for prime time.

Cognitive Systems and the PAL Program

In the 2000s, DARPA ushered in a rebirth of AI and pursued a significant thrust into cognitive systems. The flagship program in this effort was the PAL, a massive R&D effort that funded research in machine learning, natural language processing, and human-computer interaction. The main goal of PAL was to integrate the best of these techniques to create first-of-a-kind cognitive assistants that would emulate an executive assistant. Functionality included being able to understand the user's situation, manage their email, make calendar appointments, prepare for meetings, interact with the user, understand the user's tasks, anticipate next steps, and provide the user with helpful advice.

Two integration teams, one led by SRI International (Cognitive Agent that Learns and Organizes [CALO]) and one led initially by Carnegie Mellon University

DARPA Changed My Career and the Trajectory of Apple

DARPA changed my career and the trajectory of Apple, helping propel it to become the most valued company in the world.

Have you heard of Apple's Siri? Of course you have! This ambitious effort to create the first conversational assistant on your smart phone led the way for the current wave of personal assistants including Amazon's Alexa, Microsoft's Cortana, Samsung's Bixby, and Google's Google Assistant.

This would not have come about without DARPA. In 2003, DARPA launched the PAL effort, which included the CALO project at SRI International. CALO was a significant AI development effort involving more than 400 researchers from more than 25 universities across the land. The goal of the project was to build an integrated intelligent assistant that would learn "in the wild" — rather than being preprogrammed with specific knowledge on how to perform a fixed set of actions, CALO would learn every day to help its user in an office environment through observation, interaction, and reasoning.

With the support of DARPA, the CALO project made contributions on multiple levels. At the core technology level, researchers worked on advancing the state-of-the-art in machine learning and its combination with probabilistic logic; DARPA-funded publications and code contributed to the fields of statistical relational learning, Markov logic networks, probabilistic relational models, conditional random fields, latent Dirichlet allocation, Bayesian inference, and more. At the architecture level, integration architects worked to organize technical contributions into a cognitive architecture that would allow learning to drive significant parts of the system. All of this emerged in the form of an assistant called CALO, accessible through a "semantic desktop" called IRIS. (IRIS stood for "Integrate. Relate. Infer. Share" — hmmm, can you find any relation to the name Siri?)

CALO could help you with many things during your workday. About to attend a meeting? CALO would gather all of the documents, emails, and task content you needed to have at your fingertips. Working on a presentation? CALO could suggest the next slide to add, or even put together the draft of the presentation for you on any subject. Once you arrived at the meeting, you could ask CALO to record a transcription and track all task-related status and commitment updates, helping you keep your project plan up to date. CALO could also automate repetitive or mundane tasks that come up during your workday.

While the full CALO system was used daily by some members of the project, lighter-weight versions were created to help with widespread deployment, including "CALO Express," specifically built for Microsoft office users, and Active Ontologies, a drag-and-drop integration framework and AI toolkit. When SRI International decided to spin out a commercial venture, Active Ontologies was used as the technical foundation of the startup called "Siri, Inc." A few weeks after Siri was launched as a free app in the Appstore, Steve Jobs called and convinced the founders to sell their company and join Apple. With further commercialization efforts, the team launched Siri as part of the Apple iPhone 4S. Sales from that device broke all records of the time, and Apple's stock price nearly doubled, passing Exxon to make Apple the most valued company on the stock market.

It is hard to overstate DARPA's contributions. While the initial CALO assistant was a far cry from the today's personal assistants, those assistants would not exist without the sustained vision and leadership of DARPA. Decades of investment by DARPA in knowledge representation and reasoning, speech and language understanding, machine learning, and large integration projects like CALO and PAL laid the foundation for today's rapid progress in AI.

And I hope DARPA will change your life, and the trajectories of generations to come.

– Adam Cheyer

and later by SRI (Reflective Agents with Distributed Adaptive Reasoning - RADAR), pulled these ideas together into integrated prototype personal assistants. Both of these large research teams had a three-tiered organization (figure 7). The first was a basic research tier with individual research efforts in knowledge representation and reasoning, machine learning, perception, language understanding, planning and task management, and multimodal dialog. Because of the need to integrate learning and reasoning, much of the research focused on techniques for probabilities and logic (for example, statistical relational learning, Markov logic networks, probabilistic relational models, conditional random fields, latent Dirichlet allocation, and Bayesian inference). The second tier integrated the best of the component techniques into an integrated personal assistant. Tier three was an application tier which, in close collaboration with DARPA, created versions of the assistant to support military commanders.

The integration efforts (see figure 8) required the teams to select, modify, and integrate the best of the diverse research components into an integrated personal assistant that could pass annual Go/No-Go tests that measured how much the performance of the integrated system improved after a week-long learning period. This arduous task was often frustrating and sometimes counterproductive, but provided stresses that led to new insights. The chief architect for the SRI team, Adam Cheyer, not only created an integrated system, but personally used it every day, implementing his own double-helix development strategy. This eventually led to the creation of a lightweight version of the CALO system, CALO Express, and the development of a lightweight integration environment, Active Ontologies (Guzzoni, Baur, and Cheyer 2007). The full SRI CALO system provided three major sets of capabilities: information management, task management, and meeting understanding. As the DARPA effort wound down, SRI invested in turning this environment into a prototype personal assistant for the newly developed Apple iPhone and created the Siri venture that was eventually sold to Apple. As with other transitions, the capabilities developed during the DARPA program went well beyond its commercial realization.

Because of DARPA's continued support of projects reflecting man-machine symbiosis, the state-of-the-art has advanced to the point where a significant marketplace for intelligent virtual assistants is emerging and projected to be approximately $20 billion by 2025 (Hernandez 2018). The first generation represented by Siri, Alexa, Google Assistant, and Cortana have set a high bar and integrated assistants are becoming the normal expectation in many walks of life.

Conclusion

Throughout its 60-year history, DARPA has played a unique role among institutions in a position to inspire and lead advanced technological efforts in the United States: it has consistently pushed the community to tackle extremely challenging, affectionately termed *DARPA-hard*, problems whose solutions could have profoundly important outcomes; it has been adept at organizing large-scale, multi-organizational efforts; and many of its programs are consistently mindful of the big picture, concentrating on integration of capabilities and their real-world deployment in a way that is generally not seen in other agencies that sponsor research. These characteristics have made DARPA particularly well-suited to lead the research and development community in pushing the boundaries of AI, and it has been doing this with expertise since the earliest days of the field.

Importantly, as DARPA has consistently insisted, broad-based, robust AI requires the fundamental integration of a variety of capabilities to give an autonomous system the ability to sense and act on the world, reason about its circumstances and what to do next, communicate with other agents, and learn continuously from its environment. As we have outlined, the ultimate success of deployed AI systems in open, unpredictable worlds, in which circumstances change and new missions arise, will depend on the ability of perceptual, reasoning, and learning capabilities to work very closely together. DARPA realized this back in the 1960s, and has continued to drive research forward toward that very challenging goal ever since; most efforts have concentrated on large-scale integrated AI efforts on autonomous mobile platforms and intellectual partners for humans. Through a multidecade sweep that includes mobile robots, autonomous vehicles, partners for pilots, and personalized assistants, DARPA has been a primary driver, uniquely situated to help the community make important progress toward integrated, robust AI.

Acknowledgments

This project was supported by the Defense Advanced Research Projects Agency.

References

Banks, S. B., and Lizza, C. S. 1991. Pilot's Associate: A Cooperative, Knowledge-Based System Application. *IEEE Expert* 6(3): 18–29. doi.org/10.1109/64.87681

Bush, V. 1945. As We May Think. *Atlantic Monthly* 176(1): 101–8.

Chapman, R. J.; Graham, J. M.; Carley, K. M.; Rosoff, A.; and Paterson, R. 2005. Providing Insight into Command Post Operations Through Sharable Contextualized Net-Centric Visualizations and Analysis. In Proceedings of the 2005 Human Interaction with Complex Systems Symposium. Berlin, Germany: Springer.

Engelbart, D. C. 1962. Augmenting Human Intellect: A Conceptual Framework. Summary Report AFOSR-3233. Menlo Park, CA: Stanford Research Institute. doi.org/10.21236/AD0289565

Engineering and Technology History Wiki 2017. Milestones: SHAKEY: The World's First Mobile Intelligent Robot, 1972. Engineering and Technology History Wiki, February 16. ethw.org/Milestones:SHAKEY:_The_World%E2%80%99s_First_Mobile_Intelligent_Robot,_1972.

Gat, E.; Bonnasso, R. P.; and Murphy, R. 1998. On Three-Layer Architectures. In Artificial Intelligence and Mobile Robots, Kortenkamp, D., Bonasson, R. P., and Murphy, R., editors. 195–210. Cambridge, MA: AAAI Press/The MIT Press.

Grosz, B. J. 1974. The Structure of Task-Oriented Dialogs. In Proceedings of the IEEE Symposium on Speech Recognition, 250–254. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Guzzoni, D.; Baur, C.; and Cheyer, A. 2007. Modeling Human-Agent Interaction with Active Ontologies. Paper presented at the Association for the Advancement of Artificial Intelligence (AAAI) Spring Symposium on Interaction Challenges for Intelligent Assistants, March 2007. www.aaai.org/Papers/Symposia/Spring/2007/SS-07-04/SS07-04-009.pdf.

Hart, P. E., Nilsson, N. J., and Raphael, B. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. Institute of Electrical and Electronics Engineers (IEEE) Transactions on Systems Science and Cybernetics 4(2): 100–107. www.datamation.com/applications/ai-assistants-to-lend-1-billion-users-a-virtual-hand-by-2025.html.

Hernandez, P. 2018. AI Assistants to Lend 1 Billion Users a Virtual Hand by 2025. Datamation, January 2. www.datamation.com/author/21491910/Pedro-Hernandez.

Kotsuruba, I., and Tsotsos, J. K. 2018. A Review of 40 Years in Cognitive Architecture Research: Core Cognitive Abilities and Practical Applications. arXiv Preprint. arXiv:1610.08602v3 arxiv.org/pdf/1610.08602.pdf.

Laird, J. E.; Lebiere, C.; and Rosenbloom, P. S. 2017. A Standard Model of the Mind: Toward a Common Computational Framework Across Artificial Intelligence. Cognitive Science, Neuroscience, and Robotics. *AI Magazine* 38(4): 13–26. doi.org/10.1609/aimag.v38i4.2744

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

Nilsson, N.; Deutsch, B.; Fikes, R.; Sacerdoti, E.; and Tenenbaum, J. 1974. *Plan for a Computer-Based Consultant System*, SRI International Technical Note 94. Menlo Park, CA: SRI Artificial Intelligence Center.

Sacerdoti, E. 1980. Problem-Solving Tactics. *AI Magazine* 2(1): 7–15.

Thrun, S.; Montemerlo, M.; Dahlkamp, H.; Stavens, D.; Aron, A.; Diebel, J.; Fong, P.; Gale, J.; Halpenny, M.; Hoffmann, G., et al. 2006. Stanley: The Robot that Won the DARPA Grand Challenge. *Journal of Field Robotics* 23(9): 661–92. doi.org/10.1002/rob.20147

Ron Brachman is the Director of the Jacobs Technion-Cornell Institute at Cornell Tech in New York City and a professor of Computer Science at Cornell University. Previously he was the chief scientist at Yahoo and head of Yahoo Labs. He was the director of the IPTO at DARPA from 2002 to 2005. His distinguished research career focused on knowledge representation and reasoning; he coauthored with Hector Levesque a prominent textbook in that area and he is a cofounder of the International Conferences on Knowledge Representation and Reasoning. Brachman served as President of the Association for the Advancement of Artificial Intelligence from 2003 to 2005. He has won numerous awards and is a Fellow of the Association for the Advancement of Artificial Intelligence, the American Association for the Advancement of Science, the Association for Computing Machinery, and the Institute of Electrical and Electronics Engineers.

David Gunning has managed AI projects at DARPA over the last three decades as a program manager in the 1990s (1994–2000), the 2000s (2003–2008), and the 2010s (2015–2019). His programs included the CPOF in the 1990s, PAL in the 2000s, and Explainable AI in the 2010s. He recently left DARPA to become a Technical Program Manager at Facebook AI Research. Gunning holds an MS in Computer Science from Stanford University, and an MS in Cognitive Psychology from the University of Dayton.

Murray Burke was the program manager for DARPA's High Performance Knowledge Bases, Rapid Knowledge Formation, and DARPA Agent Markup Language programs. After completing his appointment at DARPA, Burke has been providing systems engineering and technical assistant services to numerous DARPA AI programs over the past 16 years, including PAL and Deep Exploration and Filtering of Text programs. Burke is currently assisting on the Explainable AI and Machine Common Sense initiatives.