DARPATech 2004
Beyond Science Fiction: Building a Real Cognitive Assistant

Mr. David Gunning
Information Processing Technology Office
Defense Advanced Research Projects Agency

For decades, Hollywood has speculated about what an intelligent computer assistant would be like. In the 1960's, it was Star Trek's talking computer and, later, it was HAL 9000. In the 70's, it was R2D2. In the 80's, it was Star Trek's Commander DATA. Today, movies like Artificial Intelligence depict entire civilizations of ultra-intelligent machines. It seems that each decade, these imaginary beings become more capable and intelligent. But they all share a common property: none of them really exist!

Now, DARPA is about to build the world's first real cognitive assistant. It won't be anything like what Hollywood has imagined, but it could have a major impact on future military operations.

Let's look at a future military commander. Here is an artist's conception of a ground forces commander in the year 2015. In fact, these images were also created in Hollywood—or at least very near by—at USC's Institute for Creative Technologies. The commander is leading a light, well equipped, objective force—with a new generation of vehicles, remote sensors, precise weapons, and, most importantly, as few people as possible.

But, like the Hollywood version of any story, these images leave out some important details. The real information environment will be intense. These are just a few of the information systems the commander will enjoy. This will be his version of the "monster" Ron Brachman wants us to tame. How else will our commander cope with this mess?

So, here is the first gap we need to bridge—the gap between the commander and the information "monster" we are likely to leave him. How do we bridge this gap today?

Today's commanders deal with a relatively meager information environment by having lots, and lots, and lots of assistants. Here is a typical command center—where hundreds of information systems come together and are staffed by hundreds or thousands of individuals. Who or what will give our future commander the assistance he now gets from hundreds of staff? In the Hollywood version of this story, we would expect the commander to be assisted by an amiable and competent computer assistant. This is the image we've just had created to convey our idea of a cognitive commander's assistant. Of course he doesn't really exist either. And, what does exist is not even close to what we need.

We have all seen computer assistants that happily pop-up on our screen, but rarely know what we need. We have all called automated help systems on the phone and been frustrated by the brittleness of their predetermined script. Interacting with one of these automated assistants is nothing like talking to Commander DATA. It is more like talking to an idiot savant from Mars.

The essential problem is that these computerized assistants are brittle—they have no ability to reason or learn. They are not cognitive.

What do we mean by cognitive? Let me walk through a simple example. Let's say you are successful here at DARPATech and are invited to visit DARPA to tell us about your new idea! But, it's your the first visit to DAPRA and you need to find your way to 3701 North Fairfax Drive. You, as a cognitive being, can easily handle this, even if you've never been to Virginia. You start with some knowledge—you know about maps, roads, traffic, and parking lots. You can reason with that knowledge—you can devise a route, estimate times, and plan for parking. You can ask questions—even if only as a last resort—you can, if absolutely necessary, ask someone for directions. You can perceive the environment—you can see where traffic is building up and adjust your route accordingly. Finally, you can learn.

Let's assume that your visit to DARPA is successful, and you win a contract. You start coming back, again, and again, and again. As a result: 1) you may realize that you should be more careful about what you wish for; and 2) each time you come back you will learn to improve your route.

You will learn to adapt to rush hour traffic. You will learn new routes from new hotels. Eventually, you will be able to make the trip in your sleep and won't need to be cognitive at all. Humans do this all the time and in every aspect of cognition. This property is well understood by psychologists and called Cognitive Skill Acquisition.

The first time we encounter a new situation, we struggle through it, reasoning, slowly if necessary, until we finally figure out a solution. The second and third time through, our performance improves dramatically until we hit the physical limits of task performance. This learning curve captures the essence of what we need our cognitive systems to do: to reason through new experiences the first time and, then, learn to improve their performance every time after that.

In contrast, imagine you had a GPS navigation system in your car. The first time through, it would give you a much better route than the one you fumbled through on your own. But, following that, it would give you exactly that same route. It would not learn. It would not learn about new traffic patterns. It would not learn new short cuts. It would not learn that you always stop at Starbucks on the way to your meeting. Worse yet, it would be limited to solving this one very specific problem.

If, on the following day, you decide to take the train; your navigation computer would not have a clue what to do. It would be totally lost, once outside of its carefully programmed domain. This is the fundamental problem we need to solve. This is the difference between a cognitive system and an idiot savant.

Building an assistant with this essential capability will not be easy. We believe it can only be done by the building of a complete cognitive system, through the fundamental integration of several new technologies.

Imagine our commander now has an automated assistant. He instructs his assistant to alert him if there is any sign of enemy activity on his left flank. How will his assistant handle that request?

First the system will have knowledge—not a database of simple facts—but knowledge about the world, such as what is meant by a flank, what are common activities a military unit performs, who is the commander, and what are his preferences.

The system will use reasoning. It will use a variety of inference, planning, and scheduling techniques. It will use reasoning to infer the exact location of the left flank in this case and calculate which areas should be monitored for enemy activity.

The system will use dialogue. It will be able to ask the commander questions—but only as a last resort. If absolutely necessary, it might ask the commander what he means by activity—does he mean any vehicle movement or just coordinated maneuvers by an entire unit.

The system will use perception. It should be able to perceive obvious events. It would know when the commander is talking and not choose that moment to ask an annoying question.

Finally, the system will learn. It will use learning in every component. In the same way that you learned to improve your route to DARPA with each trip, the commander's assistant will learn to improve with each interaction. It will add new knowledge with each experience. It will improve its reasoning with each outcome. It will improve its dialogue with each conversation. Learning is central to making the system cognitive.

It is the reason IPTO is also initiating the Real World Learning program described by Barbara Yoon—to explore as many ideas as possible in this crucial area. The focus of my effort here is to build a complete cognitive system—one that includes learning—but one that also includes reasoning, dialogue, and perception.

How do we build this system? Luckily, much of the work has been done for us. For the past twenty years, researchers have made slow but steady progress in all of the required technical areas. We have now had twenty years of "science projects" in the fields of: knowledge representation, reasoning and inference, planning and scheduling, machine learning, machine vision, speech and language understanding, multimodal dialogue, and cognitive science, to name a few.

Many of these ideas have gone well beyond science projects and become real, working systems. For example, we now have machine learning techniques that can learn to recognize credit card fraud—better than a human. We have machine learning systems that have learned to play backgammon—without any human programming—and beaten the world champion. We have reasoning systems that control deep space exploration, including diagnosing spacecraft failures and responding to unexpected events. This week we will see autonomous vehicles that are able to plan routes around obstacles in the California desert. Speech understanding systems are now commonplace. Moreover, we now have advanced dialogue systems that can carry on more natural, extended, and flexible conversations.

This explosion in new capabilities has created the final gap we need to bridge—the gap between these individual research results—these science projects—and the creation of a complete, working, cognitive system. In short, we need to bridge the gap between science and science fiction.

DARPA has just started a new program called PAL which stands for: Personalized Assistant that Learns, to do just that.

Two teams, SRI International and Carnegie Mellon University, have been awarded contracts to develop and integrate the first instances of a complete cognitive assistant. The SRI team has team members from across the country—it has both outstanding researchers in all of the required technical areas and a large integration team to bring them together. The SRI team includes coordinated research efforts in: knowledge and reasoning, dialogue, perception, and learning. These research centers will start with the 20 years of science projects and create the components needed for our cognitive agent. SRI will then integrate those components into a working prototype. Every year for the next five years, this team will produce a working cognitive assistant.

The CMU team is also bringing together the outstanding researchers in the needed technical areas. However, the CMU team happens to be bringing together the best researchers to be found on the CMU campus. They will develop and bring together their technologies into their own integrated architecture. Each year, for the next five years, CMU will produce a working prototype of their cognitive assistant. Each year, these systems will be tested to see how well they are meeting the goal of becoming a reasoning and learning cognitive assistant.

However, this is a very large and difficult problem. The two PAL teams will—I have no doubt—create something that has never existed before. Nevertheless, there are still many "science projects" yet to be explored.

A few areas come to mind that are not being fully covered by the current efforts. Here are five topics that might be worth exploring and on which we need your help.

First, Core knowledge. Any intelligent system, even one that learns, must start with a critical mass of core knowledge. What should be included in this core? How should it be represented? How can this large amount of knowledge be acquired?

Second, "a-logical" reasoning. The PAL project is using and developing a variety of logical and probabilistic reasoning techniques. However, people are not merely logical. Much of our reasoning—and often our best reasoning—is through analogy and metaphor, or by thinking in specific examples. How do we implement these "a-logical" reasoning techniques in a machine?

Third, one of the critical open challenges for these systems is the ability to reflect on their own experience and perform "meta-reasoning" as it is called. This is especially true when it comes to handling novelty and surprise. How does a system know that it is in a new situation? How does it know what it doesn't know?

Fourth, visual communication. Especially as we transition PAL into military systems, we will need to communicate visually with the commander—by allowing him to sketch his ideas and by showing him interesting visualizations of complex battlefield dynamics.

Fifth, novel cognitive architectures. SRI and CMU will explore two versions of a cognitive architecture, but surely there is room to explore other architectural ideas.

In conclusion, bring us your best ideas. We are eager to find promising new ideas that we have not heard before. With your help, we hope to fill in the missing science, bridge the gap between science and science fiction, and build—for the first time—a real cognitive assistant—and all in time to help our commander in 2015.

Thank you.