DARPATech 2004
Taming the Monster: The Next Computing Revolution
Dr. Ronald Brachman
Director
Information Processing Technology Office
Defense Advanced Research Projects Agency

Virtually everything we do in modern society rests on computers.

Think about it—from dialing your cell phone, to getting money from an ATM, to using your anti-lock brakes, not to mention actually using your PC—almost every ingredient of your day is dependent on computing technology. And the same certainly goes for our military, a situation made even clearer when we look at each of the services' visions of the future.

As you heard this morning, from FCS to ForceNet to the JBI, our future is "network-centric"—which means that at the bottom, it all depends on computers.

But there's a huge irony here. There is no doubt that without IPTO's computing inventions of the last 40 years, our new military visions would not even be imaginable.

But, along with our successes, we've created a monster. Think about your own PC. It crashes. It gives you bizarre error messages. It just goes catatonic every once in a while and the best it can do is try to keep you up to date with that marvelously communicative hourglass cursor.

And remarkably, even though you do the same things with it day after day after day, it doesn't even notice! If a bug caused it to crash yesterday, the same bug will cause it to crash again today and yet again tomorrow. It doesn't try to take advantage of its own past experience or adapt to your needs. It doesn't even know you exist. In fact, in any real sense of the word, it just doesn't know anything.

And not only are today's computing systems creatures of ignorance, they are extraordinarily complex, they take deep expertise to debug, and they are on the verge of being completely incomprehensible to human beings. They have an unending supply of vulnerabilities and failure modes. And the size of the staff that our military needs just to keep these machines running is nothing short of outrageous—and it keeps getting worse.

If we are going to survive—let alone realize the visions of our military services—we must make a radical change in our view of computing. With your help, we are going to drive this change.

The mission of DARPA's Information Processing Technology Office is to tame the monster of our own creation—by developing a new kind of computing system that can truly know what it's doing. If our nation's future is going to be so dependent on computers, we need to find a way to give them some new capabilities.

Our computer systems must be more able to look out for themselves, so that we don't need scores of humans to set them up and keep them running. They need to be able to learn about their environments over time so that once they survive an attack or a breakdown, they can immunize themselves against future occurrences. They need to be able to adapt to their users so that humans can do their jobs more effectively, and they need to be sensitive to those users' missions and priorities. And we can no longer tolerate their grinding to a halt when faced with unanticipated inputs—they must do something reasonable, even if they encounter a new situation for which they have not been explicitly programmed.

When you think through this list of essential capabilities, something interesting becomes apparent: to control the computational monster of complexity and vulnerability, we need computers to exhibit abilities that we usually think of as "cognitive."

Cognition is generally defined as having to do with knowledge and the ability to use that knowledge to make judgments, and more generally, to reason. A cognitive computing system would not only need to reason with knowledge, but it would require semantically rich ways to remember its experiences and recall them exactly when needed. It would need its input sensors to do much more than simply transmit streams of raw data—those sensors would need to more closely resemble the active kind of perceptual system that natural entities rely on. A cognitive system would also need to be able to communicate, and to accept naturally expressed guidance from others. And of course, no matter how clever its reasoning was, it couldn't stay static. Systems that don't learn over time at the very least appear stupid, and at worst they simply don't survive.

These elements—reasoning, learning, perception, remembering, and using experience, and communicating naturally—are the key ingredients of our office's strategic direction. And of course, these elements need to be integrated into a coherent whole, since it is only when they come together that we can get a truly intelligent system.

So, imagine for a moment, if you will, a cognitive network, which really understood what it was doing: it could recognize precursors of a denial of service attack based on past experience and on high-level advice from network administrators, and proactively set up defenses to avoid downtime; or an intelligent radio: it could know enough about past experience with poor reception areas to warn you before your Humvee entered a zone where you'd be likely to lose your signal, and advise you to finish your transmission before you lost connectivity; or a thoughtful command and control system: it could know what the operator on the last shift was working on and make sure that the new operator did not waste any time getting started and did not make any false assumptions, thereby avoiding potential for serious mistakes. And so on—you get the idea.

If you think about it, these new types of applications have capabilities that are reminiscent of those of good human assistants. Human assistants adapt to their superior's style and preferences, they understand the mission and act proactively to support it, they take guidance responsibly, and they use common sense to plan and re-plan on a moment-to-moment basis.

Over the last year, we've started working on ideas for what we have called an "enduring-personalized-cognitive-assistant" as a microcosm of our grand challenge of cognitive computing.

Exactly what capabilities does a proactive, thoughtful, commonsensical human assistant usually have? If you want to think of a role model, think of Radar O'Reilly, the company clerk on the old TV show, M*A*S*H. Radar was no rocket scientist, but he sure could anticipate what the Colonel needed to do before the Colonel himself realized it. He always managed to do the right thing, using his wits and general knowledge, and especially by being sensitive to the needs of his boss. Now that's the kind of system we want to build!

[Video] Thank you Radar. Visions of this sort are nice but the key, practical question is, "how do we get there from here?"

Let me tell you briefly about how we structure IPTO's programs. This will give you a good idea of the specific, pragmatic steps we are taking and exactly where you fit in. First, we have a set of core technology programs. For example, we might develop a technology thrust in what we call "real-world" learning—how to create learning algorithms that emulate the scope and effectiveness of the many ways that people really learn. Barbara Yoon will be exploring this direction with you in a few moments.

There are many other core technology ideas that we'd like to pursue this way and we are starting to think about new programs for some of them. The technology ideas include, for example, reasoning at real-world scale; natural language and multimodal communication; the creation and use of episodic memory so computers can remember and refer to their past experience; self-aware reflection; the integration of perception and cognition; and others that we need your help in fleshing out.

Second, we need programs that drive us to put these capabilities together in a tightly integrated fashion. For example, for all of the great success in machine learning over the last 15 years, the main results have been in isolated learning in narrow domains—nothing like the kind of learning that we do every day, which takes into account what we already know, advice we receive from experts, thoughtful reflection, purposeful experimentation, and other inputs. That kind of learning needs to work tightly with reasoning and perception.

This is one of our bigger challenges, but it is absolutely essential to push very hard on the integration of cognitive capabilities into unified, competent fully-functional systems. We take the "systems" side of cognitive systems very seriously.

One key issue here is, what's an appropriate underlying architecture to use to integrate learning, reasoning, perception, and action? We are looking for good ideas for system architectures that will allow tight integration between these elements yet allow for the kind of flexibility and robustness we see in natural systems. So far, I've implied that we have focused on building a single system. And in a few minutes Dave Gunning will tell you about an important new project that is attempting to build an individual cognitive system. But we need to go beyond single-agent systems.

A critical ingredient of a successful military is team performance, and teams of cognitive systems should be able to do substantially more than individual ones can. And don't forget, teams of the future will include both robots and humans. As a result, cognitive systems will need to be able to be supervised by humans and will need to act responsibly and in trustworthy fashion. Our office plans to address important issues in this general area, which we've been calling "collaborative cognition."

Here's a third possible approach to cognition: imagine a system made up of a large number of elements, each of which is very simple—for example, a swarm of bees. Overall, the collective can exhibit quite robust and intelligent behavior, if it's organized in the right way. Somehow we get global intelligence from locally limited elements.

If you're familiar with Orson Scott Card's "Ender" books, you'll have a sense of the possibilities I'm thinking about here. In that science fiction quartet, Card explored some of the possible characteristics of an alien race structured around a "hive mind." Collective cognition in computers is a promising new area, and Chris Ramming will shortly tell you about some upcoming prospects in systems that might do learning and reasoning in a distributed fashion.

There's one final, absolutely critical element to our office strategy: IPTO's new cognitive agenda will need to partner with its more traditional work in computer science. While I've focused on the cognitive part here, some of our most important programs are about computing in its more conventional form.

The really big opportunity we see is in bridging the gap between the cognitive and computing sides of our office. This is how we'll put the computing monster in its place.

There are two things to do here. First, we can focus some of our core technology development in hardware and software on novel architectures and techniques for making the cognitive algorithms most effective. Bob Graybill is starting a new program in this direction, which we call Architectures for Cognitive Information Processing, and the BAA for this program is still open.

Second, we can use novel developments in cognitive information processing technology to improve more conventional computing systems. We have high hopes for this approach. We believe it can help provide better system security, more rapid system configuration, self-diagnosis and maintenance, and increased functionality over time through learning. Lee Badger will share some of his thoughts on new program ideas in this area in a few moments.

This is really the ultimate goal of IPTO: to show how breakthroughs in cognitive computing will enable new classes of computing machines that can take care of themselves and improve over time, can act responsibly with supervised autonomy, and that can realistically support the military's vision of a computer-permeated future.

We have the beginnings of our new cognitive systems agenda in place, and we've already started a number of new programs, as you can see here. But we've just barely scratched the surface. I've talked about our general framework, but we need your help in working out the rest of the big picture, and in making it concrete.

To get into more specifics on what we need, I've invited our newest program managers and our relatively new deputy director to tell you where we're thinking of going over the next year or two. They will be eager to hear about your new ideas.

Before wrapping up, let me list some key technical challenges that our vision of cognitive computing demands we address. As I've mentioned, a cognitive system must have the ability to reason. After all, the very definition of cognition involves having knowledge and using it to reason.

There's a long tradition of work in artificial intelligence that has explored different forms of reasoning, and formal reasoning systems can do some very sophisticated things. But what still seems to be elusive is the ability to reason very pragmatically based mainly on personal experience and plausible assumptions, the way most of us seem to reason most of the time.

I believe that we need to explore some new directions in representation and reasoning that are heavily experience-based. In general, we need to find some new methods that allow a cognitive system to come to plausible, qualitative conclusions, even if they're not totally provable or quantified.

On a different front, if we are going to think of our cognitive computers as assistants, we almost certainly need to take into account the role of emotion and affect in their operation. There's surprising evidence in natural systems that emotion and decision-making go hand in hand.

Despite entertaining fictional explorations of androids without emotions like Commander Data, our current totally robotic computers always seem out of sorts, or even depressed. How to build affect into a computer and how to use it to good effect are important open questions.

As subsequent talks will emphasize, learning is one of the most important issues we face. As I've said before, systems that are incapable of learning are invariably perceived to be stupid. In order to succeed, we'll need systems that can remember where they've been and what they've seen, and improve themselves over time. But learning doesn't occur in a vacuum, and so we need to explore deeply the connection between learning and other aspects of being cognitive. And there does not appear to be only one kind of learning. We need to understand far better how to make systems learn by reading, by experimenting, by practicing, and by observation.

Finally, evaluation and metrics are critical to our cognitive enterprise. You may be aware of the "Turing Test," originally proposed by Alan Turing more than 50 years ago in which he imagined a scenario that would test a machine's ability to fool a questioner into thinking it was a human.

In some respects, this is still the only semi-serious attempt to frame the evaluation problem for complex, integrated, multi-faceted cognitive systems. But Turing's so-called "Imitation Game" is more a parlor game about deceiving an interrogator than anything else. It is overly simplistic, and totally inadequate for our needs.

Over the next few years, our office will need to develop a better understanding of how to test cognitive systems. If we can create thorough evaluation methodologies that are quantifiable, unambiguous, and can be applied year over year to confirm real progress, then perhaps we can avoid some of the problems that have plagued research like this in the past. But this is something we really don't know how to do and for which we eagerly seek your ideas.

More generally, we need your help in a related area: one of the best ways to create a new research program is to offer a vivid and thorough description of the grand challenge test to be taken at its end.

In the case of our PAL program, we have envisioned a system that performs well in specific scenarios that are exactly like those that a human executive assistant would face.

You know all about the DARPA Grand Challenge—the problem to be solved there is clear, simple, compelling, and very exciting. We urgently need your help in identifying the clear, simple, compelling and exciting grand challenges in cognitive systems. What key tests in perception, reasoning, learning, and other areas do computing systems need to pass in order to show the world that cognitive technology really matters and really works?

It's very easy to imagine how powerful a technology cognitive information processing might be. If we succeed, even if it takes us many years, we will finally be able to defeat the monster that stands between us and a computer-permeated future. We're talking here about a revolution in computing that's comparable in its impact to that of machine automation itself.

But the challenge is daunting. In many key areas, we hardly know where to begin. And some parts of this road have been traveled before, with mixed success.

But let us not forget, IPTO has changed the world more than once before. Thanks to our incredible predecessors in the original IPTO and its offspring, our office has created an amazing legacy for the modern military and for 21st century civilization as a whole.

The litany of accomplishments is as familiar as it is remarkable: time-sharing, personal computing, the ARPANet, parallel processing, RISC processors, the Internet, speech recognition, and even the creation of the field of computer science itself.

As IPTO launches itself into its fifth decade, we ask you for your help in adding one more revolutionary item to our proud history of world-changing accomplishments: computing systems that truly know what they're doing.

Thank you very much.