DARPATech 2004
Get Smart: Real World Learning

Dr. Barbara Yoon
Deputy Director
Information Processing Technology Office
Defense Advanced Research Projects Agency

Learning is a key element of any intelligent system. It's the fundamental way to acquire and assimilate new information to increase our knowledge of the world.

Why is this important? Well, for one thing, without learning, even the most seemingly intelligent entity is doomed to repeat the same mistakes endlessly.

The world is a dynamic place, but static, non-learning systems have no mechanism for dealing with unanticipated change. Humans are not born with an innate ability to function in society; rather, we learn all our adult competencies. Learning provides us the ability to deal with new situations, and ultimately, to improve our responses and behavior over time.

So, it should come as no surprise that as the Department of Defense pushes towards increased autonomy and decreased manpower in its deployed systems, learning is critical to our new military paradigms. Without learning, we would have to anticipate and engineer, in advance, appropriate responses to every conceivable contingency.

Even if this were possible, and it's not, this approach would take enormous time and resources, and would have unacceptable life cycle costs. Learning is the only effective way to produce robust systems that can survive in the field. But learning alone is not sufficient.

As Ron Brachman has said, IPTO is in the midst of a renaissance. We have embarked on a mission to create a new generation of cognitive systems.

Part of our office vision is rooted in the conviction that the highly capable cognitive systems of the future will require system integration as a fundamental scientific concept. We believe that revolutionary advances in machine learning are now going to be driven by its interplay with other components of intelligence, particularly reasoning and knowledge.

This is the theme we are exploring in a set of studies that will lay the groundwork for a future program in what we call Real World Learning. We are looking for major breakthroughs that will push machine learning beyond its current capabilities to the much broader goal of knowledge acquisition. And we envision that the most promising technologies will transition to systems such as PAL, the Personalized Assistant that Learns, which you'll hear about in a moment, and to many other applications.

Historically, machine learning has focused mainly on small-scale problems in well-defined, carefully constrained domains. The most well-developed technology is supervised learning for pattern classification and simple control tasks. This approach has yielded a rich set of applications that are in widespread use, including simple adaptive filters, credit card fraud detection, and autonomous navigation in simple environments. However, supervised learning systems need massive amounts of labeled training data and stringent training procedures to ensure proper generalization. In addition, these systems don't scale well to large-sized problems, and typically need to forget what they know before they can learn something new.

So, it's apparent that we are still far from being able to learn in a "real-world" environment, that is, in an application that is large, complex, dynamic, and fraught with competing or conflicting goals.

What do we need to get there? Let's look at a hypothetical example and review a few themes and observations about how humans learn.

Take the case of a senior medical resident learning clinical procedures. Let's say you visit the clinic complaining of a fever, a dry cough, and pain in your knee. The resident will interpret your complaint based on your medical history, his medical knowledge, and today's newspaper report of an early flu season. After the resident examines you, he concludes that your respiratory symptoms are consistent with the flu, and that your knee pain is an arthritic flare up that's unrelated. But when he finds out you've had a flu shot this year, he starts to look for other explanations. On a hunch, he asks you if you've been around birds lately. When you say yes, he first asks more questions, then consults a text on infectious diseases, and finally makes a tentative diagnosis of parrot fever, which is not a type of flu, but which has similar symptoms. His diagnosis is seconded by the instructing doctor and eventually confirmed by follow-up tests. From this example, first, observe that, as Ron said, learning never occurs in a vacuum. Rather, we always learn from an existing body of knowledge.

Context strongly influences what you look for, what new information you retain, and how you assimilate it into what you already know. Furthermore, the more you know, the faster and deeper you learn. So, the resident, unlike a layman, drew on his wealth of medical knowledge to recognize meaningful patterns of information, to focus on the right questions to ask, and to learn a broader interpretation of flu-like symptoms, which by the way, doesn't include pain in the knee. The resident's education is much more than just a large body of facts.

Medical knowledge—or for that matter, any expert knowledge—is organized around core concepts into a schema of inter-related models. Every new piece of information, including your symptoms, is used to form a current "model of the world" that helps the resident efficiently retrieve relevant information, and transfer learned knowledge from one situation to another.

In the National Academy of Sciences study on "How People Learn", evidence is presented that indicates when experts retrieve from memory information that's relevant to a particular problem, they don't exhaustively search through everything they know about the problem. Instead, they selectively and effortlessly retrieve only the appropriate information. They're able to do this because the knowledge they acquire includes a specification of the contexts for which it's useful. You might think of this as a tag. So, in effect, they can go fluently right to the "answers".

Studies have also shown that learning conceptual abstractions and invoking analogy enhance a person's ability to transfer knowledge from one situation to another, and that interestingly, developing a suite of representations for a complex problem domain enables people to think flexibly about that and related domains. Needless to say, these notions haven't yet been incorporated into machine learning technology. So we need your help in addressing this challenge.

A second thing to notice is, when we learn, we do so at many levels, some of which are not even conscious. In this case, the medical resident had an intuition—let's say, based on a subconscious memory from childhood having nothing to do with his medical education—and that led him to ask you about birds. Intuition can be a powerful force in guiding decisions and problem solving. But it isn't usually a conscious process.

In his book, "An Anthropologist on Mars", the neurologist, Oliver Sacks, describes a patient who is institutionalized with a type of brain damage that prevents him from learning anything new. The patient could not consciously remember where he was or who he had met, and yet he was able to learn at a subconscious level how to navigate around the institution and who and what situations were pleasant or unpleasant.

There is clearly a hierarchy of learning levels. But how do these different levels interact to produce robust learning? Here again, we look to you for help.

A third point to consider, learning is driven by reasoning. Learned information is always being considered and reconsidered for its appropriate place in the "scheme of things". When you told the resident that you had knee pain, he had to think for a second to create and reject several hypotheses and explore the possible implications of your statement. Eventually he concluded that it didn't fit with your other symptoms, and was instead a second, unrelated ailment. Furthermore, at some point, he realized that he needed to consult a reference text on infectious diseases to confirm his hypothesis. In essence, reasoning is the key to orchestrating learning.

In its broadest sense, reasoning, taken to its full potential, allows you to deduce your strengths and weaknesses, decide what you need to learn and how best to do it, and organize new information into transferable knowledge.

Fundamentally, learning should involve acquiring new knowledge about the world, and new knowledge should enable new reasoning, which will refine and enhance learning. This is a very ambitious goal, one that lies at the heart of the long-standing allure of artificial intelligence, and one that most critically needs new creative ideas from you.

These are some of the issues that have led us to consider the framework illustrated here for our new learning thrust, one that places learning algorithms in the broader context of reasoning and knowledge structures. This framework is a simplified schematic of how learning fits into a larger picture.

A few things to note are: Currently the main areas are largely separate technologies with separate technical communities. So, to say it again, the message here is integration and interaction.

Second, representation is a key issue. It's not clear whether there is a single best representation for a particular problem or body of information, or whether multiple representations are needed.

An even thornier issue is the integration of fundamentally mismatched representations in an extensible way, such as semantic and statistical representations.

Third, we are interested in new multi-faceted learning techniques that can learn robustly from sparse data, and seamlessly switch among a variety of modes, including direct advice, reading, focused inquiry and experimentation, one-shot learning from a single experience, and other modes.

And finally, as Ron has said, perhaps the key issue here is, what is the appropriate architecture for this integrated system concept?

To better understand how to address these challenges, we have initiated a set of studies that are exploring several new ideas in real-world learning. As I've said previously, these studies will help lay the groundwork for a future program in this area.

Let me now touch briefly on a few of them. The first problem area could be called "learning in the large": how to handle huge domains that approach the complexity of the real world.

The goal here is to draw from clues on how humans get around in the world with such seeming ease. Granted, adults have a very large base of knowledge, both specialized and technical as well as so-called "common sense" knowledge. But it's also the case that, when you have to decide which way to swerve to avoid an oncoming car, you don't freeze for minutes on end to "compute" a rigorous, "perfect" answer. Instead, you produce a more intuitive quick response. You don't have to find the optimum solution; you only need one that keeps you alive and unharmed.

Consider now that you have to make decisions like this every day and in many different areas of expertise. Your life experience includes proposals, differential equations, gourmet cooking, poetry, dancing, high-rise offices, and so on. It's impossible to consider action selection over a domain this big. Instead, it's more likely that we solve huge problems by breaking them down into a sequence of smaller sub-problems.

Our knowledge and goals are hierarchical in nature, and we are constantly using this multi-resolution hierarchy of abstractions to formulate and solve "bite-sized" problems.

We don't yet know how to design artificial systems to do this.

Although our study is addressing some promising approaches, there are still a number of really difficult problems. For example, how do you learn an enormous model of the world in tractable pieces and integrate them? How do you compose your tractable problem from diverse pieces? How do you learn and plan in parallel at different levels of resolution, from short-term specific to long-term abstract? Where do our hierarchies come from? And so on.

This last question brings us to the second problem area: developing hierarchies of abstractions and concepts from specific examples and low-level processes. This is a key issue in forming knowledge from learned data. Why is this important?

Suppose you learn about distance, speed, and travel time while planning a car trip, but can't generalize that knowledge to an airplane trip. What good does it do to learn something if you can't use it in new situations? So knowledge transfer is, in part, a measure of the quality of acquired knowledge.

Humans seem to rely strongly on analogy. We frequently invoke metaphor. One of our studies is exploring how to extract general principles and concepts from analogies, that is, multiple similar instances. But this is a very hard problem. How do you ensure that a system will extract useful concepts rather than silly, spurious ones?

Another study in this area is exploring techniques for deriving high-level "models of how the world works" from learned low-level procedures. Observations on how people learn to make travel plans indicate that they learn more effectively by actually doing the work rather than by being told an abstract framework of procedures. Domain knowledge is built up by performing this task repeatedly under varying circumstance.

These exploratory studies are just the first step in developing a program aimed at broadening machine learning into knowledge acquisition and beyond. It has been said that learning techniques will be the basis of every application that involves a connection to the real world."

To be useful in the military world (and other settings), a single system will need to be fieldable in multiple environments and capable of "lifelong" learning.

I've described some of our exploratory studies in the hopes of stimulating your interest and creativity, but the efforts I've outlined are by no means the only approaches to this difficult problem.

As we go forward to a new program in Real World Learning, we look to you for those far-reaching ideas that will take learning to new and revolutionary dimensions of capability.

Thank you.