Beyond the Turing Test

Itamar Arel and Scott Livingston
University of Tennessee

Creating an artificial general intelligence roadmap could help researchers realize AI's original goal.

What is intelligence? Can it be objectively measured? More importantly, can it be constructed artificially?

AI researchers have been struggling with these fundamental questions for decades. The grand challenge, of course, involves much more than building adaptive word processors or clever spam filters; it is to engineer broadly competent intelligence eventually matching that of humans. Researchers refer to this as *artificial general intelligence* (AGI) to distinguish it from modern connotations of AI.

AI's original goal was as ambitious as the name of the field suggests: the realization of human-level intelligence in machines. However, in response to the great difficulties they encountered, researchers narrowed their focus, leading to the emergence of computer vision, natural-language processing, machine learning and planning, and other AI subfields.

Great advances in these areas have indeed been made—almost every new gadget seems to incorporate some adaptive behavior or awareness of usage patterns. However, a generally intelligent system, comparable in ability even to cats or dogs, has yet to be built.

RESURRECTING AI

The potential impact of a system that captures individual AI successes and combines them into something greater and holistic is enormous. Such an effort would be analogous to the invention of machines that met, and now greatly exceed, humans' physical abilities. Extending the automation and power in skyscrapers and airplanes to the realm of cognition is exciting indeed. Creating a broadly competent and intelligent robot would certainly drastically change the quality of modern life.

Advances in various AI disciplines have led to a wealth of theories and demonstrated applications that may motivate and support efforts to achieve AGI. However, numerous significant challenges remain, and even what general path to take is the subject of heated debate among researchers. For example, do we need to understand intelligent behavior in biological systems before constructing it artificially?

Moreover, there is no comprehensive, convincing, and well-supported account of intelligence itself; formulating such a theory will likely require major advances in neuroscience and cognitive science.

AGI ROADMAP

There is widespread interest in returning to AI's original goal. The First International Conference on AGI convened in 2008, and the next will take place in March 2009 (www.agi-09.org). The focus of this event, and of the *Journal of Artificial General Intelligence* (http://journal.agi-network.org) born out of it, is creating and analyzing human-level artificial intelligence.

Nonetheless, conceptions of AGI among researchers differ: What will human-level intelligence embodied in a machine look like? What will it be able to do, and what are good metrics for measuring its performance?

The famous Turing test, which requires a machine to fool a person into thinking it is also human by way of a text-based correspondence, is clearly insufficient; there are already chat-bots that can mimic a person, at least in the context of an Internet chatroom. A generally intelligent system must not only have wide applicability but also be able to face entirely new challenges.

The growing community of researchers interested in AGI must create a roadmap. Such a map should define basic notions including the type of behavior an AGI system should exhibit. It would focus the energies of

 Published by the IEEE Computer Society

those working toward this common goal while precisely clarifying intermediate targets.

A roadmap also encourages active external verification of research procedures and results by posing benchmark tests or concrete progressive steps. A recent example of a similar effort is the Virtual Worlds Roadmap (www.virtualworldsroadmap.org), which attempts to accelerate the development of massively multiplayer online worlds.

Of course, calling for a roadmap is much easier than actually making one, and undoubtedly there will be much initial disagreement.

The first step is to determine the fundamental characteristics of "true" artificial general intelligence. For example, many important problems such as automobile traffic optimization are modeled as partially observable Markov decision processes. Any generally intelligent system must be able to solve a wide range of POMDPs without making significant assumptions about the underlying process or environment with which the system interacts.

Perhaps more important, any AGI system must be able to capture both short- and long-term temporal and spatial dependencies, a skill that humans readily employ but that thus far is beyond even the most sophisticated AI platforms. This may require reformulating current optimal control methods or even proposing new economic utility modeling schemes.

GENOMIC ENDOWMENTS AND THE AI CHALLENGE

Yet another question researchers must address is the difference between learned AGI behaviors and those prescribed by the technological equivalent of *genomic endowments*—innate behavioral patterns explicitly included in the system design. A classic example is the human ability to recognize a face's basic structural features.

Such AI endowments should be limited—perhaps to the extreme of a tabula rasa. This would prevent repeating historical mistakes by essentially creating a look-up table for problem-domain-specific solutions. As the precise architecture of a successful AGI system is unknown, no final claim can be made regarding the embodied or situated nature of a promising design.

A key aspect of a useful AGI roadmap will be clearly defined steps toward the ultimate goal of human-level intelligence.

A key aspect of a useful AGI roadmap will be clearly defined steps toward the ultimate goal of human-level intelligence. Well-recognized problems should serve as part of increasingly difficult milestones. Solving each of these challenges will be a substantial achievement in itself and will move the state of the art forward.

However, an important pitfall to avoid is introducing problems that narrow AI techniques can solve. Such benchmarks would fail to distinguish general or strong AI systems from existing narrow or weak ones.

As the RoboCup competition (www.robocup.org) shows, posing several complex problems could stimulate and unify research aimed at constructing broadly competent and generally intelligent systems. Informally, it would create a common stage for sharing and proving the value of AGI research.

Despite the technological marvel of the Internet and the rapidly proliferating mobile technologies that are fundamentally changing the way we interact, AI's original "grand dream" remains elusive as we approach the twilight of the first decade in the 21st century. By designing a roadmap to AGI and creating important benchmarks, we may yet achieve that dream. However, this will only happen if the nascent AGI community coalesces and works toward a common vision.

For those interested in discussing and contributing to the construction of an AGI roadmap, a new forum can be found at agiroadmap.org.

Itamar Arel *is an assistant professor in the Electrical Engineering and Computer Science Department at the University of Tennessee, Knoxville. Contact him at itamar@ieee.org.*

Scott Livingston *is an undergraduate research assistant in the Machine Intelligence Lab at the University of Tennessee, Knoxville. Contact him at scl@utk.edu.*

Editor: Richard G. Mathieu, Dept. of Computer Information Systems and Management Science, College of Business, James Madison Univ., Harrisonburg, VA; mathierg@jmu.edu