Theoretical contributions of AI

Subbarao Kambhampati
Arizona State University

Theoretical AI is primarily the study of algorithms and representations that support efficient, approximate, continual, and resource-bounded reasoning. This significantly broadens the traditional computer science emphasis on correctness upon termination and worst-case space and time complexities.

The computational tasks arising in AI are critically linked to the types of representation languages in which reasoning is done. Mathematical logic has been the language of choice for representing declarative knowledge. Researchers have investigated a variety of logic subclasses that aid reasoning by limiting expressiveness. The area of deductive databases is a crossroads, where researchers from AI and databases consider how to represent and efficiently reason about logical knowledge.

The inherent uncertainty in most real-world knowledge has also spurred work on default logics and probabilistic representations, both of which support the incremental modifications of conclusions as evidence accumulates. Default (nonmonotonic) logics have been hailed as the first major extension of mathematical logic in a long time. Bayesian networks, which can compactly represent the dependency relationships among a set of random variables, have become the de facto representation standard for propositional uncertain knowledge. Although inference with these networks is still intractable in the worst case, many efficient approximation algorithms have been developed. Developing representations for first-order probabilistic knowledge is an active area of current research.

At the heart of AI enterprise lies the problem of helping an autonomous agent decide what to do next. Planning and reasoning about actions have thus driven many of the theoretical advances. Several representational restrictions for modeling action under uncertain and incomplete knowledge have been developed. Reasoning itself has moved from explicit simulations on world states to the more efficient action-centered representations that support direct manipulation of partially specified courses of action. This, in turn, allows plan synthesis through an efficient search of potential courses of action, with the sets being refined (narrowed) through goal-directed reasoning. The normative basis for action selection, traditionally used for decision making.

Knowledge-based systems are the most visible contribution of AI to industrial applications. They apply a simple theoretical idea: Symbolic reasoning guided by heuristics over declaratively specified knowledge of a domain can result in impressive problem-solving ability.

GRAND CHALLENGES OF AI

A Grand Challenge is a seemingly reasonable problem that is exciting and challenging yet currently unsolvable. Solutions to AI's Grand Challenges will require major new insights and fundamental advances in computer science and artificial intelligence and, if they're successful, can be expected to have a major impact on society.

Each of these tasks requires long-term, stable funding at significant levels. Success is by no means guaranteed, and each problem represents a high-risk, high-payoff investment. However, even partial success can have spin-offs to industry and have a major impact on competitiveness.

TRANSLATING TELEPHONE. A translating telephone is a system in which a Japanese speaker can converse with, say, an English speaker in real time. This requires solutions to a number of currently unsolved problems: a speech recognition system capable of recognizing a large (possibly unlimited) vocabulary and spontaneous, unrehearsed, continuous speech; a natural-sounding speech synthesis preserving speaker characteristics; and a natural language translation system capable of dealing with ambiguity, nongrammaticality, and incomplete phrases.

ACCIDENT-AVOIDING CAR. An accident-avoiding car

1982 Columbia Data Products produces the first IBM PC "clone." Compaq soon follows with its own version.

1982 Autodesk is founded and ships the first version of AutoCAD later that year.

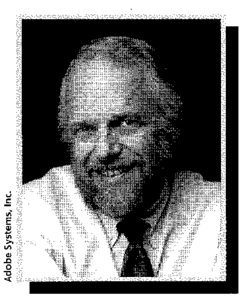

Adobe Systems, Inc.

1982 John Warnock develops the PostScript page-description language and with Charles Geschke founds Adobe Systems.

1982 Time magazine names the computer as its "Man of the Year."

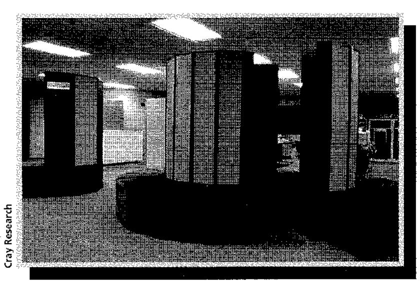

Cray Research

1982 The Cray X-MP (two Cray-1 computers linked in parallel) proves three times faster than a Cray-1.

1982

provided by Bayesian decision theory, has been extended to consider the costs and benefits of deliberation.

Computational tasks considered in AI are complicated by the fact that the agent may have severe limitations on the computational resources available to it. Resource limitations are handled by a new breed of *any-space* or *any-time* algorithms. A variety of search algorithms have been developed that can work with arbitrarily low memories and guarantee optimal solutions upon termination. Although arbitrary time restrictions cannot be met given the mostly intractable computations, a reasonable requirement is that the algorithms in question be flexible (anytime). This means they can be interrupted at any time and they return results whose value monotonically increases with increased resources. Such anytime algorithms simplify the problem of deciding how many resources to expend on decision-making and how much on execution because they offer a regular relationship between computation time and value.

Many computational tasks arising in AI admit only approximate solutions. A canonical example is inductive learning, which requires inferring a function by looking at a finite set of its values. PAC (probably approximately correct) theory provides a normative basis for evaluating approximate algorithms. Here, an algorithm's performance is measured in terms of the likelihood that it makes less than a certain percentage of output errors. In addition to providing a formal foundation for machine learning, the PAC framework started the fertile field of *computational learning theory*, in which AI and computer science theorists look at the complexity of inductive tasks.

Given the worst-case intractability of many AI problems, researchers have naturally focused on understanding the properties of "average problems." An important recent insight in this area is that many worst-case intractable decision problems have a very narrow region of transition. In other words, they are easily decided in a positive way on one side of the region and easily decided in a negative way on the other side. The hard instances lie mostly in the transition area. An important open issue is the relation between these hard instances and the instances that arise naturally in the real world.

In summary, theoretical AI provides both an increasingly solid foundation for the development of intelligent autonomous agents and a medium of rich interactions with several disciplines of computer science.

Subbarao Kambhampati *is an associate professor at Arizona State University, where he directs the Yochan group on planning and machine learning. Contact him at rao@asu.edu; http://rakaposhi.eas.asu.edu/yochan.html.*

equipped with an intelligent cruise control using sonar, laser, and vision sensors could eliminate 80 percent to 90 percent of the fatal accidents and cost less than 10 percent of the automobile's total cost. Such a device would require research in vision, sensor fusion, obstacle detection and avoidance, low cost/high speed (over a billion operations per second) digital signal processor chips, and the underlying software and algorithm design.

LEARNING SYSTEMS. Interest has long focused on systems that learn and discover from examples, observations, and books. Two long-term grand challenges for systems that acquire capability through learning are to read a chapter in a college freshman text (say, physics or accounting) and answer the questions at the end of the chapter, and learn to assemble an appliance (such as a food processor) from observing a person doing the same task. Both are extremely hard problems requiring advances in vision, language, problem-solving techniques, and learning theory. Both are essential to the demonstration of a self-organizing system that acquires capability through (possibly unsupervised) development.

SELF-REPLICATING SYSTEMS. What capabilities must exist for a factory to make a copy of itself? Self-replicating systems are of some practical interest in areas such as manufacturing in space. Rather than uplifting a whole factory, is it possible to have a small set of machine tools that can

1982 Japan launches its "fifth generation" computer project, focusing on artificial intelligence.

1982 Commercial e-mail service begins among 25 cities.

1982 In November, Compaq unveils an IBM-compatible portable PC.

Lotus

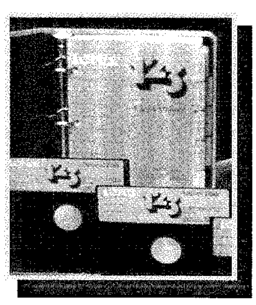

1983 By including graphics such as pie charts and bar graphs, Lotus 1-2-3 does for the IBM PC what VisiCalc did for the Apple II.

1983 A Josephson junction is developed on the basis of Brian Josephson's 1962 prediction, bringing higher speed and lower power dissipation to ICs.

1983 The IBM PC-XT heads for market success, while the PC Junior faces quick extinction.

1983 Completion of the TCP/IP switchover marks the creation of the global Internet.

1982 — 1983