New Horizons in Commercial and Industrial AI

Toshinori Munakata
Guest Editor

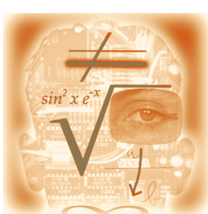

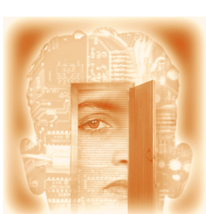

This month's Communications *presents the second of two special issues on practical applications of artificial intelligence (AI). The first issue published in March 1994 primarily focused on the most practical areas of AI. This issue presents more futuristic areas. I believe these two issues serve as milestones in presenting the current state and future perspective of everyday commercial and industrial applications of AI.*

AI as a field has undergone rapid growth in diversification and practicality. For the past 10 years, the repertoire of AI techniques has evolved and expanded. Scores of newer fields have recently been added to the traditional domains of practical AI. Although much practical AI is still best characterized as advanced computing rather than *intelligence*, applications in everyday commercial and industrial settings have certainly increased, especially since 1990. Additionally, AI has shown a growing influence on other computer science areas, such as databases, software engineering, distributed computing, computer graphics, user interfaces, and simulation.

With help from experts, I have selected promising or rapidly emerging areas of AI. The issue is divided into the following three sections:

Large-scale AI and the Commonsense Problem

The commonsense problem has been a key problem since the beginning of AI, and it is likely to continue to be so well into the next century. The problem is essential not only to AI, but also to many areas of computer science and future human civilization. For example, people understand what they see or hear by integrating with complex background knowledge, which we call common sense. This has been built over years of perceptual experience. The current AI technology, whether for a perception problem, such as machine vision or speech understanding, or knowledge processing such as expert systems or natural language processing, does not incorporate common sense. Robotics is considered one of the most successful domains within AI having practical applications. However, when we look more closely at its level of sophistication, it is similar to other areas of AI. Symbolic machine learning, despite its fairly long history, has had a relatively small number of practical successes. The relation between the commonsense problem and machine learning is a chicken-or-egg situation; any major advancement of one area may depend on the other. Similarly, any development in the common sense problem will definitely affect all areas in AI. A major breakthrough, if any, will come through an integration of commonsense with the current AI technology.

There has been much controversy on how to proceed with research on the commonsense problem. The basic problem is that we do not know how commonsense is developed and exercised in a human or even an animal brain. We tend to think commonsense is a type of knowledge that a 10-year-old would have, but a mere aggregation of pieces of information may not solve the problem. Many AI researchers believe a large-scale system will be necessary, although probably not sufficient. Current systems for the commonsense problems are based on traditional symbolic AI; typically they are extensions of either knowledge based or natural language processing systems. Non-symbolic AI, such as neural networks, appear to be too premature at present for the problem. A recent trend appears to place more emphasis on installing specific cases and corpora rather than on trying to derive everything from basic axioms or models. The latter approach is not practical at present, while the former strategy requires much larger space. Also, the history of physics, chemistry, and other sciences may suggest there might be simple rules beneath the seemingly complex commonsense problem.

The first three articles in this special issue describe ambitious projects related to the problem. We have an interesting forum here, where each author also critiques the other two projects, then rebuts the critiques of his own work. The first article by Lenat gives an introduction to the CYC project. Although there has been controversy in the CYC case, it has been a pioneer system in the world to attack the problem in the scope and current scale of about 1 million rules (reduced from 1.5 million by compaction). It is then followed by an article for the WordNet by Miller, where he jokes about his project being called "a poor man's CYC." The EDR electronic dictionary by Yokoi follows, describing a $100 million project in Japan involving huge Japanese/English language data and knowledge bases. Although all three projects have natural language orientation in various flavors, their aims are to tackle the commonsense problem.

The fourth article by Chaib-draas discusses a topic closer to realization—industrial applications in distributed AI (DAI). DAI is a distributed system whose components are AI, or intelligent agents. Through interactions among the agents, we aim at improving the performance of each individual agent, as well as the system as a whole. Since concurrent processing and information networking, such as the "information super-highway," are definite future trends in general, DAI has a strong potential for massive applications.

Advances in Traditional AI

Machine learning is a well known classical AI areas. Its basic idea—the machine gets smarter by itself—is a fascinating concept that every AI researcher dreams of. For example, give a problem to your PC. It then replies that the problem is too difficult to be solved immediately using the current knowledge base. You leave the machine on over night. It keeps working hard, deducing more conclusions and inducing or discovering new rules, internally, as well as through a possible outside interaction, such as retrieving necessary information from external knowledge and databases. When you come back the next day, the machine has become more knowledgeable than the day before and can solve the problem.

Although the number of practical applications of machine learning has been relatively small in comparison with other areas, such as expert systems, it is one of the fastest growing areas of AI. Once it is successful, its technical and social impact will be extremely significant. As in other areas of practical AI, there is a possibility that truly profitable systems are being kept secret.

The machine learning area begins with an article by Langley and Nobelist Simon, who offer a brief introduction to the area, then delves into techniques of traditional machine learning and fielded applications of these techniques. In particular, the article considers rule induction and its recent applications in greater detail, and draws strategies and lessons from successful applications. The article by Bratko and Muggleton complements its predecessor by discussing recent advances in inductive logic programming (ILP), a special form of symbolic machine learning based on the first-order predicate logic. Logic has played an important role in AI, and ILP is an application of logic to machine learning.

The second area covered in this section is natural language processing (NLP). Commercialization of many aspects of NLP has, of course, been well established. The areas of NLP range from relatively simple aids to word processing (such as spelling or elementary grammar checking), and assistance with information management on various forms of text, to much harder machine translation. A rule of thumb is that the simpler the application, the more practical.

Church and Rau survey the current state of the field.

Machine translation can demonstrate how one can make use of practical applications out of difficult areas of AI. As is common in many other challenging areas of AI, the current level of machine translation is well below that of the human translator; the exact level depends on the degree of difficulty of the target sentences. Assume, for example, that a machine translation can achieve 80% of the human level. Pessimists may emphasize negative aspects and ignore the practical use of such technology, while optimists may try to develop a system that achieves full human capability. However, the odds for a near future success for full machine translation capability are very low, and a more realistic and productive strategy is to take the middle ground. Initially, let the machine translate the text and get the 80% result, then let the human improve it to 100%. By this two-step human-machine cooperation, we may save, say, 30% of the total time for a pure human translation. There are also cases where this type of two-step technique may not work well. In a worst case scenario, the machine can mangle the text so badly that it takes the human longer to fix the damage than translate from scratch. However, by using great care when selecting the target problems and deploying machine translation, one can benefit from the technology.

Born in the context of AI, especially from such concepts as frames and inheritance, object-oriented (OO) technology has become a major computer science domain encompassing such areas as application programming, databases, operating systems, communication networks, among others. Harmon overviews OO AI—a hybrid field of OO and primarily symbolic AI—including knowledge-based systems. Although the current commercialization of OO AI is not widespread, considering the popularity of OO in general, it is one of the most promising areas for extensive implementation in the near future.

Emerging Technologies

In this section, we selected newer, somewhat exotic areas for practical applications. Although their foundations are not necessarily very recent, most commercial and industrial applications of these technologies started after 1990.

Rough sets, or approximation sets, deviate from the idea of ordinary sets. Likewise fuzzy sets vary from ordinary sets. The area is relatively new, and has remained unknown to most of the computing community. The technique is particularly suited to reasoning about vague and uncertain data, and discovering relationships in such data. It is commonly compared to other techniques, including statistical analysis, particularly discriminant analysis. Rough sets might do a better job than statistical analysis for certain problems, especially when the underlying distribution deviates significantly from a normal distribution or perhaps when the sample size is small. Pawlak, the father of rough sets, and Grzymala-Busse, Slowinski, and Ziarko introduce the principles and recent practical applications of rough sets.

The next two articles discuss chaotic systems, or simply chaos. Ditto and Munakata offer an overview of the area, while Aihara and Katayama introduces the state of chaos engineering in Japan. Chaos refers to nonlinear deterministic systems that exhibit sustained irregularity and extreme sensitivity to initial conditions. Because of their complexity, chaos systems have been shunned by most of the scientific community, despite their commonness. Recently, however, there has been growing interest in the practical applications of these systems.

Artificial life is a relatively new discipline that covers a broad spectrum of fields. Artificial life studies or creates man-made life, an approach opposite to that of biology where natural life is studied. For example, during the three-billion year history of the earth, biological life was born, perhaps as a result of a series of rare chance chemical and physical reactions of molecules. One day, artificial life might simulate such a process within a computer. Artificial life also includes such areas as evolutionary computing, artificial bacteria, simulated ecology, insect-like robots, and so forth. After a brief introduction to artificial life in general, Maes discusses a relatively practical specific topic in the area—applications of intelligent agents to the entertainment industry.

Conclusion

The term AI was coined in 1956 and since that time has come to suggest many things to many people. Indeed, as Executive Editor Diane Crawford pointed out in the first special issue, you ask five computing professionals to define "AI" and you're likely to get five different answers. AI is broadly defined as "the science of making computers do things that the human needs intelligence to do." Many people relate their definition of AI to the appearance of the problems to be solved. A computer adding 2 + 2 and giving 4 is not intelligent; a computer performing symbolic integration of $sin^2 x\, e^{-x}$ is intelligent. Classes of problems requiring intelligence have included inference based on knowledge, reasoning with uncertain or incomplete information, various forms of perception and learning, and applications to such problems as control, prediction, classification, and optimization. However, AI is also related to the simulated biological processes used to arrive at a solution, such as the use of neural networks or genetic algorithms. This view of AI is important even if such techniques are used to compute things that do not appear intelligent.

Hype related to the promise of AI, especially by the media and public, seemed to peak in the late 1980s. At that time, AI was proclaimed to be applicable to scores of problems, while its actual practical use was relatively scarce. However, since about 1990, AI has been used in increasingly practical and widespread applications. It appears some of the real gains in AI are occurring after the over-enthusiasm. Despite sometimes exaggerated views about its practicality, AI will continue to grow in practical application.