Preface

The practical side of the field of artificial intelligence (AI) aims to imbue smartness into computers and applications so that people may use the computers as if the machines are able savants. However, computers and computer systems, connected through the Internet, have evolved into giant and complex man-made machinery serving virtually anyone, anywhere and anytime. Although this Internet-connected machinery serves millions of people daily, there are many more potential users unable to use the technology because of nonintuitive and hard-to-use interfaces, limited and inflexible capabilities, delays, and outages. As our dependence on computers increases, there is an increasing need for these computers to be made smart, and the enabling technologies will depend more and more on advances in AI.

A smart machine should (1) understand the tasks to be performed, (2) perform the required tasks, and (3) manage itself with the goal of staying available and performing efficiently. These requirements translate directly into goals for the practical side of AI. In fact, the first two goals have been the main driving force for practical AI research to date. The first goal is being addressed by progress in speech and handwriting recognition, as well as advances in dialog management and natural language processing, advances that led to more intuitive and better interfaces. For the second goal, results of research in traditional areas of AI such as inferencing, knowledge representation, machine learning, searching, and constraint satisfaction have been making their way into various applications.

The third goal involves automation leading to smart computing platforms and infrastructure, and it is at the heart of the recent IBM initiative on autonomic computing. Although emphasis on this aspect of AI is new, the goal goes back to the early days of cybernetics, when Ross Ashby likened intelligence to the ability to keep the organism's essential variables constant against external disturbances. The organism of interest is the computer system and the essential variables are certain system states that are to be maintained within some acceptable range against the changes in environment. Achieving autonomic capabilities will require progress in AI and is today an additional compelling reason to pursue research in AI.

This special issue of the *IBM Systems Journal* contains thirteen papers that relate to all three goals mentioned above. They indicate the extent to which the techniques of AI are being used today and how they may be applied in the future. A Technical Forum section contains two contributions that discuss recent AI-related workshops held at the IBM Thomas J. Watson Research Center.

Sowa's "Architectures for intelligent systems" describes the organizational principles for an architecture for intelligent systems. The architecture proposed is flexible in order to optimize service to different user groups: linguists, application developers, and end users. Sowa proposes a modular framework and an AI glue language, rich enough to include first-order logic and metalevels that can state whatever speech act is intended. Two such languages are Sowa's own Conceptual Graphs and the Knowledge Interchange Format, and these can also be translated to and from some controlled natural language, to be used as a human-computer interface.

In "ABLE: A toolkit for building multiagent autonomic systems," Bigus et al. describe a Java**-based agent-building framework called ABLE. Components

are provided for building various types of agents, such as agents that communicate, sense, make inferences on various cognitive levels, or model themselves and their environment as an autonomic system.

In their paper "Intelligent probing: A cost-effective approach to fault diagnosis in computer networks," Brodie, Rish, and Ma use a Bayesian network approach for a method for intelligent diagnostic probing of computer networks. The quality of their local-inference approximation scheme degrades gracefully under increasing uncertainty, and increases as the quality of the probe set increases.

Bin et al., the authors of "Using a constraint satisfaction formulation and solution techniques for random test program generation," generate random system verification tests based on a constraint satisfaction problem formulation. The technique has been successfully used for verifying the design of several IBM products including PowerPC* processors.

In their paper "Using fuzzy control to maximize profits in service level management," Diao, Hellerstein, and Parekh present a fuzzy control scheme intended to maximize profit in an e-commerce environment in which the revenue is based on the number of completed transactions and there are penalties for service guarantee violations. Compared to classical controllers, the fuzzy controller is easier to construct and more robust for a wide range of workloads and profit models.

The remaining eight papers deal with machine learning techniques. In "Automated generation of model cases for help-desk applications," Weiss and Apte describe a new lightweight document clustering algorithm effective in the high-dimensional feature space of the word sets used in databases of customer-service problem descriptions and their resolutions. This clustering is then used to create model cases, to help users resolve their problems through self help.

In their paper "A decision-tree-based symbolic rule induction system for text categorization," Johnson et al. present a new rule induction technique that derives text classification rules from an efficient specialized decision tree for the high dimensional feature space. The rule induction system has been incorporated in several IBM products.

"A probabilistic estimation framework for predictive modeling analytics" by Apte et al. describes the ProbE tool for generating segmented predictive models from several model families. It is an industrial-strength data mining engine designed to be extensible, embeddable, and scalable. Two key applications of the tool, one for insurance risk management and the other for advanced target marketing, are also described.

In their paper "Cross training and its application to skill mining," Oblinger et al. consider the mismatch between the training set for a classification model and the application phase examples the model applies to. A case in point is cross training, that is categorizing skills of individuals from the text of their electronic communications, and applying the model to classify the skills of another individual from the entire collection of the person's e-mail. The efficacy of the solution is demonstrated through this skill-mining problem.

In "Predictive algorithms in the management of computer systems," Vilalta et al. discuss three case studies in which predictive models were developed and used to deal with three types of potential failures in systems management. In "Discovering actionable patterns in event data," Hellerstein, Ma, and Perng deal with constructing and maintaining correlation rules used in automated real-time operations systems. Although skilled experts are currently required to identify patterns in data in order to construct correlation rules, the authors propose a method to obtain those patterns by mining historical data. These two papers, as well as the paper by Diao, Hellerstein, and Parekh on fuzzy control, are indicative of the AI work required for the design of self-monitoring and self-healing components of an autonomic system.

In "Machine learning in a multimedia document retrieval framework," Perrone, Russell, and Ziq use machine learning techniques for the retrieval of handwritten documents and show their technique to be robust to machine transcription errors. They also develop a software architecture in which machine learning algorithms for feature extraction and matching can be easily integrated. In "Applying machine learning to automated information graphics generation," Zhou, Ma, and Feng use machine learning techniques to automatically obtain concise rules for information graphics generation. Clustering is also employed to help users find examples from the graphics database relevant to their needs.

In the Technical Forum section, Brackenbury and Ravin cover the Machine Intelligence and the Tur-

ing Test Workshop of last year, and McCarthy et al. cover the Symposium on Architectures for Commonsense Computing held earlier this year. The summaries of these meetings point to future work on the applications of AI.

The next issue of the *Journal* will focus on data integration in the enterprise and real-time analytics for business intelligence.

Se June Hong, Guest Editor
Alex Birman, Associate Editor
John J. Ritsko, Editor-in-Chief

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of Sun Microsystems, Inc.