Preface

According to Torsun, the area of Distributed Artificial Intelligence (DAI) deals with "cooperative problem solving using a federation of collaborating agents." These software agents are intelligent in the sense that they are adaptive, independent, and posses reasoning capability. They can plan and execute tasks in cooperation with other agents in order to satisfy their goals. They also learn in the process and adapt to changes in the environment dynamically. These agents are often heterogeneous. A multi-agent system (MAS) is defined as a loosely coupled network of problem solvers that work together to solve problems that are beyond the individual capabilities or knowledge of each problem solver. The increasing interest in MAS research is due to significant advantages inherent in such systems, including their ability to solve problems that may be too large for a centralized single agent, provide enhanced speed and reliability, and tolerate uncertain data and knowledge. Some of the key research issues related to problem-solving activities of agents in a MAS are in the areas of coordination, negotiation, and communication.

Similar to multi-agent systems, where agents communicate and collaborate with each other to accomplish a goal, in collaborative applications, a set of programs communicate with each other using proprietary protocols to let users collaborate with each other by accessing shared information. This may be accomplished by running the same application on different machines. In such a configuration, these applications are essentially clients that accept user input to drive updates to shared state and propagate the updates to other collaborating clients. With advances in Web technologies, collaborative applications are now server based and the user interface is typically a Web browser. Thus, a collaborative application can be a Web-based solution that runs on a local server that allows people communicate and work together, share information and documents, and talk in real-time over the Internet.

Recently, much research has been conducted in distributed artificial intelligence and collaborative applications. Several interesting methodologies and systems have been developed in areas such as distributed multi-agent systems for decision support, Web search and information retrieval, information systems modeling, and supply chain management. In particular, intelligent agents and multi-agent systems, and a myriad of applications have been built in various domains. This book discusses a number of agent-based applications developed for knowledge-driven decision support, online auctions, federated information systems, and mobile computing. Similarly, in the area of search and retrieval, the book provides current research in search engine performance, user centered approach for information and image retrieval, Web mining, and document clustering. Event and information modeling is an active area of research and this book presents a few chapters highlighting the salient aspects of modeling. Finally, in the area of supply chain management, forecasting demand and supply network optimization are important areas of investigation and the book presents the latest research. The following sections briefly outline the different applications discussed in various chapters of this book.

BOOK ORGANIZATION

Section I
Distributed Agent Applications and Decision Support

This section presents seven chapters that discuss distributed agent applications and decision support. Chapter I by Hong Lin, discusses a program construction method based on γ-Calculus. The problem to be solved is specified by first-order predicate logic and a semantic verification program, which is constructed directly from the specification. This method synthesizes the architectural specifications of multi-agent systems (MAS) in γ-Calculus based on the logic specifications of the MAS. By enabling the transformation from the logic specifications to operational specifications of MAS, this method allows the design of the MAS to be focused on the architectural definition level. It benefits the development of MAS by enabling logic deduction on behaviors of the MAS, and a design methodology in an incremental fashion. This chapter presents a case study of designing a course information management system.

In Chapter II, Rahul Singh presents a "Multi-Agent Architecture for Knowledge-Driven Decision Support". Organizations use knowledge-driven systems to deliver problem-specific knowledge over Internet-based distributed platforms to decision-makers. Increasingly, artificial intelligence (AI) techniques for knowledge representation are being used to deliver knowledge-driven decision support in multiple forms. This chapter illustrates how to represent and exchange domain-specific knowledge in XML-format through intelligent agents to create exchange and use knowledge to provide intelligent decision support. It shows the integration of knowledge discovery techniques to create knowledge from organizational data; and knowledge repositories (KR) to store, manage, and use data by intelligent software agents for effective knowledge-driven decision support. Implementation details of the architecture, its business implications, and directions for further research are discussed.

Chapter III by Farid Meziane and Samia Nefti is titled "A Decision Support System for Trust Formalization," in which the authors point out that trust is widely recognized as an essential factor for the continual development of business-to-customer (B2C) electronic commerce (EC). Many trust models have been developed, however, most are subjective and do not take into account the vagueness and ambiguity of EC trust and the customers' intuitions and experience when conducting online transactions. In this chapter, the authors describe the development and implementation of a model using fuzzy reasoning to evaluate EC trust. This trust model is based on the information customers expect to find on an EC Web site and that is shown from many studies to increase customers trust towards online merchants. They argue that fuzzy logic is suitable for trust evaluation as it takes into account the uncertainties within EC data and like human relationships; it is often expressed by linguistic terms rather then numerical values. The evaluation of the proposed model is illustrated using four case studies and a comparison with two other models is conducted to emphasize the benefits of using fuzzy decision system.

Mehdi Yousifi-Monod and Violaine Prince investigate learning and communication between cognitive artificial agents in Chapter IV. They attempt to answer the question: Is it possible to find an equivalency between a communicative process and a learning process, to model and implement communication and learning as dual aspects of the same cognitive mechanism? The chapter focuses on a learning situation where two agents, in a "teacher/student" relationship, exchange information with a learning incentive (on behalf of the student), according to a socratic dialog. The teacher acts as the reliable knowledge source, and the student is an agent whose goal is to increase its knowledge base in an optimal way. This chapter first defines the nature of the addressed agents, the types of relation they maintain, and the structure and contents of their knowledge base. It emphasizes the symmetry between the interaction and knowledge management, by highlighting knowledge "repair" procedures launched through dialogic

means. The chapter describes learning goals and strategies, student and teacher roles within both dialog, and knowledge handling. It also provides solutions for problems encountered by agents. A general architecture is then established.

In Chapter V titled "Improving E-Trade Auction Volume by Consortium," Sungchul Hong, Barin Nag and Dong-qing Yao present a two-tier supply chain composed of multiple buyers and multiple suppliers. They have studied the mechanism to match trading parameter, specifically volume in this study, between buyers and suppliers. The chapter discusses the architecture of the agent, and the agent community when there is cooperative matching of volume. They present a Dynamic Programming algorithm to describe the agent's decision process and heuristic algorithms as the practical solution methodology. The results of extensive experiments show the improvement achieved by the cooperation.

Chapter VI by Manoj A. Thomas, Victoria Yoon, and Richard Redmond is titled "Extending Loosely Coupled Federated Information Systems Using Agent Technology". Different FIPA complaint agent development platforms are available for developing multi-agent systems. FIPA compliance ensures interoperability among agents across different platforms. Although most agent implementation platforms provide some form of white and yellow page functionalities to advertise and identify agent roles and descriptions, there is no clear architectural standard that defines how agent community can effortlessly adapt to operate in Federated Information System (FIS), where new content sources are constantly added or changes are made to existing content sources. This chapter presents a framework based on the Semantic Web vision to address extensibility in a loosely coupled FIS.

In Chapter VII, H. Hamidi and K. Mohammadi discuss "Modeling Fault Tolerant and Secure Mobile Agent Execution in Distributed Systems". The reliable execution of mobile agents is a very important design issue in building mobile agent systems and many fault-tolerant schemes have been proposed so far. Security is a major problem of mobile agent systems, especially when monetary transactions are concerned. Security for the partners involved is handled by encryption methods based on a public key authentication mechanism and by secret key encryption of the communication. To achieve fault tolerance for the agent system, especially for the agent transfer to a new host, they use distributed transaction processing. They propose a fault-tolerant approach for mobile agent design, which offers a user transparent fault tolerance that can be activated on request, according to the needs of the task. They also discuss how transactional agents with different types of commitment constraints can commit transactions. Furthermore, they present a solution for effective agent deployment using dynamic agent domains.

Section II
Search and Retrieval

This section presents six chapters dealing with various aspects of search engines, information retrieval, and Web mining. Xiannong Meng and Song Xing in Chapter VIII, titled "Search Engine Performance Comparisons," present a comparative analysis of Microsoft Search Engine (MSE), AllTheWeb and Yahoo. In a few comparisons, other search engines such as Google, Vivisimo are also included. The study collects statistics such as the average user response time, average process time for a query reported by MSE, as well as the number of pages relevant to a query reported by all search engines involved. The chapter also studies the quality of search results generated by MSE and other search engines using *RankPower* as the metric. They found that MSE performs well in speed and diversity of the query results, while weaker in other statistics, compared to some other leading search engines. The contribution of this chapter is to review the performance evaluation techniques for search engines and use different measures to assess and compare the quality of different search engines, especially MSE.

Antonio Picariello and Antonio Rinaldi, in Chapter IX titled "A User-Centered Approach for Information Retrieval," emphasize that information retrieval can gain great advantages and improvements by considering users' feedback. The user dimension is a relevant component that must be taken into account while planning and implementing real information retrieval systems. In this chapter they first describe several concepts related to relevance feedback methods, and propose a novel information retrieval technique which uses the relevance feedback concepts in order to improve accuracy in an ontology-based system. In particular, they combine the semantic information from a general knowledge base with statistical information using relevance feedback. Several experiments and results are presented using a test set constituted of Web pages.

In Chapter X titled "Classification and Retrieval of Images from Databases Using Rough Set Theory," Jafar Ali and Aboul Ella Hassanien present an efficient algorithm to classify and retrieve images from large databases using rough set theory. Color and texture are two well-known, low-level perceptible features to describe an image contents. The features are extracted, normalized, and then the rough set dependency rules are generated directly from the real value attribute vector. Then the rough set reduction technique is applied to find all reducts of the data which contains the minimal subset of attributes that are associated with a class label for classification. They have tested three different popular distance measures and found that quadratic distance measures provide the most accurate and perceptually relevant retrievals. The retrieval performance is measured using recall-precision measure, as is standard in all retrieval systems.

Chapter XI by Lars Werner and Stefan Böttcher is titled "Supporting Text Retrieval by Typographical Term Weighting". Text documents stored in information systems usually consist of more information than the pure concatenation of words, for instance they also contain typographic information. Because conventional text retrieval methods evaluate only the word frequency, they miss the information provided by typography (e.g., regarding the importance of certain terms). In order to overcome this weakness, the authors present an approach which uses the typographical information of text documents and shows how this improves the efficiency of text retrieval methods. Their approach uses weighting of typographic information in addition to term frequencies for separating relevant information in text documents from the noise. They have evaluated their approach on the basis of automated text classification algorithms. The results show that their weighting approach achieves very competitive classification results using at most 30% of the terms used by conventional approaches, which makes their approach significantly more efficient.

Ben Choi and Xhongmei Yao discuss a Web mining approach in Chapter XII titled "Web Mining by Automatically Organizing Web Pages into Categories". Since the majority of Web content is stored in the form of Web pages, this chapter focuses on techniques for automatically organizing Web pages into categories. Various Artificial Intelligence techniques have been used; however the most successful ones are classification and clustering. Clustering is well suited for Web mining by automatically organizing Web pages into categories each of which contains Web pages having similar contents. However, one problem in clustering is the lack of general methods to automatically determine the number of categories or clusters. For the Web domain, until now there is no such method suitable for Web page clustering. To address this problem, this chapter describes a method to discover a constant factor that characterizes the Web domain and proposes a new method for automatically determining the number of clusters in Web page datasets. This chapter also proposes a new bi-directional hierarchical clustering algorithm, which arranges individual Web pages into clusters and then arranges the clusters into larger clusters and so on until the average inter-cluster similarity approaches the constant factor. Having the constant factor together with the algorithm, this chapter provides a new clustering system suitable for mining the Web.

In Chapter XIII, John Goh, and David Taniar discuss "Mining Matrix Pattern from Mobile Users". Mobile user data mining is about extracting knowledge from raw data collected from mobile users. There have been a few approaches developed, such as frequency pattern, group pattern, parallel pattern, and location dependent mobile user data mining. Previously proposed methods share the common drawbacks of costly resources that have to be spent in identifying the location of the mobile node and constant updating of the location information. The proposed method aims to address this issue by using the location dependent approach for mobile user data mining. Matrix pattern looks at the mobile nodes from the point of view of a particular fixed location rather than constantly following the mobile node itself. This can be done by using sparse matrix to map the physical location and use the matrix itself for the rest of mining process, rather than identifying the real coordinates of the mobile users. This allows performance efficiency with slight sacrifice in accuracy. As the mobile nodes visit along the mapped physical area, the matrix will be marked and used to perform mobile user data mining. The performance and evaluation shows that the proposed method can be used for mobile user data mining.

Section III
Information Systems and Modeling

This section presents four chapters that discuss the use of modeling in information systems, particularly, event and information modeling, predictive modeling, and causal strategy modeling. In Chapter XIV titled "Conceptual Modeling of Events for Active Information Systems," Sal March and Gove Allen argue that the ontological foundations for active information systems must include constructs that represent concrete and conceptual objects, their attributes, and the events that affect them. Events are a crucial component of conceptual models that represent active information systems. The representation of events must include ascribed attributes representing data values inherent in the event as well as rules defining how conceptual and concrete objects are affected when the event occurs. The state-history of an object can then be constructed and reconstructed by the sequence of events that have affected it. Alternate state-histories can be generated based on proposed or conjectured rule modifications, enabling a reinterpretation of history. Future states can be predicted based on proposed or conjectured events and event definitions. Such a conceptualization enables a parsimonious mapping between an active information system and the organizational system in which it participates.

Chapter XV by John Artz discusses "Information Modeling and the Problem of Universals". Earlier work in the philosophical foundations of information modeling identified four key concepts in which philosophical groundwork must be further developed. This chapter reviews that earlier work and expands on one key area – the Problem of Universals – which is at the very heart of information modeling.

In Chapter XVI titled "Empirical Inference of Numerical Information into Causal Strategy Models by Means of Artificial Intelligence," Christian Hillbrand observes that many companies build their strategy upon poorly validated hypotheses about cause and effect of certain business variables. However, the soundness of these cause-and-effect-relations as well as the knowledge of the approximate shape of the functional dependencies underlying these associations turns out to be the biggest issue for the quality of the results of decision supporting procedures. It is sufficiently clear that mere correlation of time series is not suitable to prove the causality of two business concepts. However, one can find proven causality techniques in other sciences like econometrics, mechanics, neuroscience, or philosophy. This chapter presents an approach which applies a combination of well-established statistical causal proofing methods to strategy models in order to validate them. These validated causal strategy models are then used as the basis for approximating the functional form of causal dependencies by the means of Artificial

Neural Networks. This in turn is employed to build an approximate simulation or forecasting model of the strategic system.

Chapter XVII by Yongjian Fu, Hironmoy Paul, and Namita Shetty is titled "Improving Mobile Web Navigation Using N-Grams Prediction Models". This chapter discusses the use of N-gram models for improving Web navigation for mobile users. N-gram models are built from Web server logs to learn navigation patterns of mobile users. They are used as prediction models in an existing algorithm which improves mobile Web navigation by recommending shortcuts. The experiments on two real data sets show that N-gram models are as effective as other more complex models in improving mobile Web navigation.

Section IV
Supply Chain Management

This section contains two chapters that discuss supply chain management using intelligent technologies such as machine learning and petri nets. Chapter XVIII by Réal Carbonneau, Rustam Vahidov, and Kevin Laframboise is titled "Forecasting Supply Chain Demand using Machine Learning Algorithms". In this chapter, the authors point out that demand prediction in a supply chain is aggravated by the fact that communication patterns between participants tend to distort the original consumer's demand and create high levels of noise. Distortion and noise negatively impact forecast quality of the participants. This chapter investigates the applicability of Machine Learning (ML) techniques and compares their performances with the more traditional methods in order to improve demand forecast accuracy in supply chains. The authors use two data sets from particular companies (chocolate manufacturer and toner cartridge manufacturer), as well as data from the Statistics Canada manufacturing survey. A representative set of traditional and ML-based forecasting techniques have been applied to the demand data and the accuracy of the methods was compared. As a group, Machine Learning techniques outperformed traditional techniques in terms of overall average, but not in terms of overall ranking. A Support Vector Machine (SVM) trained on multiple demand series produced the most accurate forecasts.

In Chapter XIX "Supporting Demand Supply Network Optimization with Petri Nets", Teemu Tynjala proposes a generic methodology for describing and analyzing demand supply networks, (i.e. networks from a company's suppliers through to its customers). There can be many possible demand supply networks with different logistics costs for a product. Therefore, a Petri Net-based formalism, and a reachability analysis-based algorithm that finds the optimum demand supply network for a user-specified product structure is introduced. This method has been implemented and is currently in production use inside all Nokia business groups. It is used in demand supply planning of both network elements and handsets. An example of the method's application to a concrete Nokia product is discussed in the chapter.

Effective use of distributed artificial intelligence technologies and collaborative applications is critical for improving productivity. An outstanding collection of latest research associated with artificial intelligence techniques, collaborative applications, agent-based systems, Web search and mining, information systems modelling, and supply chain management is presented in this book. Use of intelligent information technologies will greatly assist in gaining competitive advantage.

Vijayan Sugumaran
Editor-in-Chief
Distributed Artificial Intelligence, Agent Technology, and Collaborative Applications