

Preface

Computational intelligence (also called artificial intelligence) is a branch of computer science that explores methods of automating behavior that can be categorized as intelligent. The formal study of topics in computational intelligence (CI) has been under way for more than 50 years. Although its intellectual roots can be traced back to Greek mythology, the modern investigation into computational intelligence can be traced back to the start of the computer era, when Alan Turing first asked if it would be possible for “machinery to show intelligent behaviour.” Modern CI has many sub-disciplines, including reasoning with uncertain or incomplete information (Bayesian reasoning, fuzzy sets, rough sets), knowledge representation (frames, scripts, conceptual graphs, connectionist approaches including neural networks), and adaptive and emergent approaches (such as evolutionary algorithms and artificial immune systems).

CI has a long history in business applications. Expert systems have been used for decision support in management, neural networks and fuzzy logic have been used in process control, a variety of techniques have been used in forecasting, and data mining has become a core component of Customer Relationship Management (CRM) in marketing. More recently developed agent-based applications have involved the use of intelligent agents — Web-based shopping advisors, modelling in organizational theory and marketing, and scenario-based planning in strategic management. Despite the obvious benefits of CI to business and industry - benefits of modeling, forecasting, process control and financial prediction to name only a few - practitioners have been slow to take up the methods available.

Business practitioners and researchers tend to read and publish in scholarly journals and conference proceedings in their own discipline areas. Consequently, they can be unaware of the range of publications exploring the interaction between business and computational intelligence. This volume addresses the need for a compact overview of the diversity of applications of CI techniques in a number of business disciplines. The volume consists of open-solicited and invited chapters written by leading international researchers in the field of business applications of computational intelligence. All papers were peer reviewed by at least two recognised reviewers. The book covers some

foundational material on computational intelligence in business, as well as technical expositions of CI techniques. The book aims to deepen understanding of the area by providing examples of the value of CI concepts and techniques to both theoretical frameworks and practical applications in business. Despite the variety of application areas and techniques, all chapters provide practical business applications.

This book reflects the diversity of the field — 43 authors from 13 countries contributed the 22 chapters. Most fields of business are covered — marketing, data mining, e-commerce, production and operations, finance, decision-making, and general management. Many of the standard techniques from computational intelligence are also covered in the following chapters — association rules, neural networks, support vector machines, evolutionary algorithms, fuzzy systems, reinforcement learning, artificial immune systems, self-organizing maps, and agent-based approaches.

The 22 chapters are categorized into the following seven sections:

Section I: Introduction

Section II: Marketing Applications

Section III: Production and Operations Applications

Section IV: Data Mining Applications

Section V: Management Applications

Section VI: Financial Applications

Section VII: Postscript

Section I contains three chapters, which provide introductory material relating to CI applications in business. Chapter I provides an overview of the field through a cross-sectional review of the literature. It provides access to the vast and scattered literature by citing reviews of many important CI techniques, including expert systems, artificial neural networks, fuzzy systems, rough sets, evolutionary algorithms, and multi-agent systems. Reviews and cited articles cover many areas in business, including finance and economics, production and operations, marketing, and management. Chapter II identifies important conceptual, cultural and technical barriers preventing the successful commercial application of CI techniques, describes the different ways in which they affect both the business user and the CI practitioner, and suggests a number of ways in which these barriers may be overcome. The chapter discusses the practical consequences for the business user of issues such as non-linearity and the extrapolation of prediction into untested ranges. The aim is to highlight to technical and business readers how their different expectations can affect the successful outcome of a CI project. The hope is that by enabling both parties to understand each other's perspective, the true potential of CI in a commercial project can be realized. Chapter III presents an innovative use of CI as a method for collecting survey-type data in management studies, designed to overcome "questionnaire fatigue." The agent-based simulation approach makes it possible to exploit the advantages of questionnaires, experimental designs, role-plays, and scenarios, gaining a synergy from a combination of methodologies. The chapter discusses and presents a behavioral simulation based on the agent-based simulation life cycle, which is supported by Web technology. An example

simulation is presented for researchers and practitioners to understand how the technique is implemented.

Section II consists of four chapters illustrating marketing applications of CI (Chapters IV to VII). Chapter IV develops a heuristic genetic algorithm for product portfolio planning. Product portfolio planning is a critical business process in which a company strives for an optimal mix of product offerings through various combinations of products and/or attribute levels. The chapter develops a practical solution method that can find near optimal solutions and can assist marketing managers in product portfolio decision-making. Chapter V reviews some classical methods for modeling customer brand choice behavior, and then discusses newly developed customer behavior models, based on boosting and stacking neural network models. The new models are applied to a scanner data set of liquid detergent purchases, and their performance is compared with previously published results. The models are then used to predict the effect of different pricing schemes upon market share. The main advantage of these new methods is a gain in the ability to predict expected market share. Chapter VI reviews several fields of research that are attempting to solve a problem of knowledge management related to the retrieval and integration of data from different electronic sources. These research fields include information gathering and multi-agent technologies. The chapter uses a specific information gathering multi-agent system called MAPWeb to build new Web agent-based systems that can be incorporated into business-to-consumer activities. The chapter shows how a multi-agent system can be redesigned using a Web-services-oriented architecture, which allows the system to utilize Web-service technologies. A sample example using tourism information is presented. Chapter VII uses a data-mining information retrieval technique to create a Web-mining system. It describes how an off-line process is used to cluster users according to their characteristics and preferences, which then enables the system to effectively provide appropriate information. The system uses a fuzzy c-means algorithm and information retrieval techniques that can be used for text categorization, clustering and information integration. The chapter describes how this system reduces the online response time in a practical test case of a service Web site selling mobile phones. The case shows how the proposed information retrieval technique leads to a query-response containing a reasonable number of mobile phones purchase suggestions that best matched a user's preferences.

Section III contains three chapters illustrating CI applications in the general field of production and operations (Chapters VIII to X). Chapter VIII discusses the various techniques, such as artificial neural networks, wavelet decomposition, support vector machines, and data mining, that can be used for the forecasting of market demand and price in a deregulated electricity market. The chapter argues that the various techniques can offer different advantages in providing satisfactory demand and price signal forecast results, depending on the specific forecasting needs. The techniques can be applied to traditional time-series-based forecasts when the market is reasonably stable, and can also be applied to the analysis of price spikes, which are less common and hence more difficult to predict. Chapter IX presents a hybrid-agent model for Belief-Desire-Intention agents that uses CI and interactive learning methods to handle multiple events and intention reconsideration. In the model, the agent has knowledge of all possible options at every state, which helps the agent to compare and switch between options quickly if the current intention is no longer valid. The model uses a

new Adaptive Neuro-Fuzzy Inference System (ANFIS) to simulate vessel berthing in container terminals. The chapter shows how the agents are used to provide autonomous decision making capabilities that lead to an enhancement of the productivity of the terminal. Chapter X describes a new CI algorithm called Horizon Scan, a heuristic-based technique designed to search for optimal solutions in non-linear space. Horizon Scan is a variant of the Hill-Climbing technique. The chapter describes an application of the technique to finding the optimal solution for the scheduling-pricing-dispatch problem in the Australian deregulated electricity market. The approach outlined is general enough to be applied to a range of optimization problems.

Section IV consists of five chapters in the general area of data mining (Chapters XI to XV). Chapter XI argues that data-mining algorithms often generate a large number of rules describing relationships in the data, but often many of the rules generated are not of practical use. The chapter presents a new technique that integrates visualization into the process of generating association rules. This enables users to apply their knowledge to the mining process and be involved in finding interesting association rules through an interactive visualization process. Chapter XII suggests using association rule data-mining techniques to assist manufacturing companies with customer requirement analysis, one of the principal factors in the process of product development. Product development is an important activity in an organization's market expansion strategy. In situations where market segments are already established and product platforms have been installed, the methodology can improve the efficiency and quality of the customer requirement analysis process by integrating information from both the customer and design viewpoints. The chapter argues that generating a product portfolio based on knowledge already available in historical data helps to maintain the integrity of existing product platforms, process platforms, and core business competencies. A case study of vibration motors for mobile phones is used to demonstrate the approach. Chapter XIII suggests that, while association rules mining is useful in discovering items that are frequently found together, rules with lower frequencies are often of more interest to the user. The chapter presents a technique for overcoming the rare-item problem by grouping association rules. The chapter proposes a method for clustering this categorical data based on the conditional probabilities of association rules for data sets with large numbers of attributes. The method uses a combination of a Kohonen Self-Organizing Map and a non-linear optimisation approach, combined with a graphical display, to provide non-technical users with a better understanding of patterns discovered in the data set.

Chapter XIV provides a brief historical background of inductive learning and pattern recognition. It then presents an introduction to Support Vector Machines, which belong to a general class of problem solving techniques known as kernel methods. The chapter includes a comparison with other approaches. As the chapter points out, the basic concept underlying Support Vector Machines is quite simple and intuitive, and involves separating out two classes of data from one another using a linear function that is the maximum possible distance from the data. While free and easy-to-use software packages are available, the actual use of the approach is often impeded by the poor results obtained by novices. The chapter aims at reducing this problem by providing a basic understanding of the theory and practice of Support Vector Machines. Chapter XV presents an overview of one of the oldest and most fundamental areas in data mining, that of association rule mining. It also introduces the maximum sub-array

problem, an approach that is gaining importance as a data-mining technique. A number of other data-mining algorithms, covering decision trees, regression trees, clustering, and text mining, are also briefly overviewed. The chapter provides pseudo-code to demonstrate the logic behind these fundamental approaches to data mining, and gives online access to code to enable CI practitioners to incorporate the algorithms into their own software development.

Section V considers management applications, particularly tools and support for decision-making, in three chapters (Chapters XVI to XVIII). Chapter XVI introduces a new deliberative process to enhance group decision-making within organizations, by allowing for and against propositions in a discussion to be explicitly articulated. The approach is called ConSULT (Consensus based on a Shared Understanding of a Leading Topic), and provides a computer-mediated framework to allow for asynchronous and anonymous argumentation, collection and evaluation of discussions, and group decision-making. The approach can be used in conjunction with any CI technique to enhance the outcome of group decision-making. Chapter VII describes an uncertain-reasoning-based technique called NCaRBS (N state Classification and Ranking Belief Simplex), an extension of the CaRBS system developed from Dempster-Shafer theory. The chapter shows how the technique can be used to categorize the strategic stance (Prospector, Defender, or Reactor) of U.S. states in relation to the public provision of long-term care. The approach also has the advantage of treating missing values, which are very common in most public sector data, as ignorant evidence rather than attempting to transform them through imputation. The system displays the results graphically, which the authors argue helps the elucidation of the uncertain reasoning-based analysis, and which should help move public management research towards better benchmarking and more useful examinations of the relationship between strategy and performance. Chapter XVIII argues that simple multi-criteria decisions are made by first deriving priorities of importance for the criteria in terms of a goal, and then priorities of the alternatives in terms of the criteria identified. Benefits, opportunities, cost and risks are also often considered in the decision-making process. The chapter shows how to derive priorities from pair-wise comparison judgments from theories of prioritisation and decision-making using the Analytic Hierarchy Process (AHP) and the Analytic Network Process (ANP), both developed by the author. The techniques are illustrated with a number of examples, including an estimation of market share.

Section VI contains three chapters demonstrating financial applications (Chapters XIX to XXI). Chapter XIX introduces artificial immune system algorithms, inspired by the workings of the natural immune system and, to date, not widely applied to business problems. The authors point out that the natural immune system can be considered as a distributed, self-organising, classification system that operates in a dynamic environment and, as such, has characteristics that make its simulated equivalent very suitable for offering solutions to business problems. The chapter provides an example of how the algorithm can be used to develop a classification system for predicting corporate failure. The chapter reports that the system displays good out-of-sample classification accuracy up to two years prior to failure. Chapter XX presents an intelligent trading system, using a hybrid genetic algorithm and reinforcement learning system that emulates trader behaviour on the Foreign Exchange market and finds the most profitable trading strategy. The chapter reports the process of training and testing on historical data, and shows that the system is capable of achieving moderate gains over the period

tested. The chapter also reports the development of real-time software capable of replacing a human trader. Chapter XXI provides an overview of recent online portfolio selection strategies for financial markets. The aim of the strategies is to choose a portfolio of stocks to hold in each trading period, using information collected from the past history of the market. The chapter presents experimental results that compare the performance of these strategies with respect to a standard sequence of historical data, and that demonstrate future potential of the algorithms for online portfolio selection. The chapter suggests that investment companies are starting to recognize the usefulness of online portfolios trading for long-term investment gains.

Finally, in Section VII, after the technical material of the preceding chapters, the post-script (Chapter XXII) presents a non-technical topic, a brief overview of the history of mathematics-based approaches to problem solving and analysis. Despite the tremendous gains in our theoretical understanding and practical use of statistics and data analysis over the last half century, the discipline remains grounded in the work of early pioneers of statistical thought. The chapter shows the human dimension of these early developments from pre-history through to the beginning of the 20th century.

This book will be useful to business academics and practitioners, as well as academics and researchers working in the computational intelligence field who are interested in the business applications of their areas of study.