Preface

Discussion on the nature of intelligence long pre-dated the development of the electronic computer, but along with that development came a renewed burst of investigation into what an artificial intelligence would be. There is still no consensus on how to define artificial intelligence: Early definitions tended to discuss the type of behaviours which we would class as intelligent, such as a mathematical theorem proving or displaying medical expertise of a high level. Certainly such tasks are signals to us that the person exhibiting such behaviours is an expert and deemed to be engaging in intelligent behaviours; however, 60 years of experience in programming computers has shown that many behaviours to which we do not ascribe intelligence actually require a great deal of skill. These behaviours tend to be ones which all normal adult humans find relatively easy, such as speech, face recognition, and everyday motion in the world. The fact that we have found it to be extremely difficult to tackle such mundane problems suggests to many scientists that an artificial intelligence cannot simply display the high-level behaviours of an expert but must, in some way, exhibit some of the low-level behaviours common to human existence.

Yet this stance does not answer the question of what constitutes an artificial intelligence but merely moves the question to what common low-level behaviours are necessary for an artificial intelligence. It seems unsatisfactory to take the stance which some do, that states that we would know one if we met one. This book takes a very pragmatic approach to the problem by tackling individual problems and seeking to use tools from the artificial intelligence community to solve these problems. The techniques that are used tend to be those which are suggested by human life, such as artificial neural networks and evolutionary algorithms. The underlying reasoning behind such technologies is that we have not created intelligences through such high-level techniques as logic programming; therefore, there must be something in the actuality of life itself which begets intelligence. For example, the study of artificial neural networks is both an engineering study in that some practitioners wish to build machines based on artificial neural networks which can solve specific problems, but it is also a study which gives us some insight into how our own intelligences are generated. Regardless of the reason given for this study, the common rationale is that there is something in the bricks and mortar of brains — the actual neurons and synapses — which is crucial to the display of intelligence. Therefore, to display intelligence, we are required to create machines which also have artificial neurons and synapses.

Similarly, the rationale behind agent programs is based on a belief that we become intelligent within our social groups. A single human raised in isolation will never be as intelligent as one who comes into daily contact with others throughout his or her developing life. Note that for this to be true, it is also required that the agent be able to learn in some way to modulate its actions and responses to those of the group. Therefore, a pre-programmed agent will not be as strong as an agent which is given the ability to dynamically change its behaviour over time. The evolutionary approach too shares this view in that the final population is not a pre-programmed solution to a problem, but rather emerges through the processes of survival-of-the fittest and their reproduction with inaccuracies.

Whether any one technology will prove to be the central one in creating artificial intelligence or whether a combination of technologies will be necessary to create an artificial intelligence is still an open question, so many scientists are experimenting with mixtures of such techniques.

In this volume, we see such questions implicitly addressed by scientists tackling specific problems which require intelligence with both individual and combinations of specific artificial intelligence techniques.

OVERVIEW OF THIS BOOK

In **Chapter I**, Tran, Abraham, and Jain investigate the use of multiple soft computing techniques such as neural networks, evolutionary algorithms, and fuzzy inference methods for creating intelligent decision support systems. Their particular emphasis is on blending these methods to provide a decision support system which is robust, can learn from the data, can handle uncertainty, and can give some response even in situations for which no prior human decisions have been made. They have carried out extensive comparative work with the various techniques on their chosen application, which is the field of tactical air combat.

In **Chapter II**, Tsoi, To, and Hagenbuchner tackle a difficult problem in text mining — automatic classification of documents using only the words in the documents. They discuss a number of rival and cooperating techniques and, in particular, give a very clear discussion on latent semantic kernels. Kernel techniques have risen to prominence recently due to the pioneering work of Vapnik. The application to text mining in developing kernels specifically for this task is one of the major achievements in this field. The comparative study on health insurance schedules makes interesting reading.

Bai and Zhang in **Chapter III** take a very strong position on what constitutes an agent: “An intelligent agent is a reactive, proactive, autonomous, and social entity”. Their chapter concentrates very strongly on the last aspect since it deals with multi-agent systems in which the relations between agents is not pre-defined nor fixed when it is learned. The problems of inter-agent communication are discussed under two headings: The first investigates how an agent may have knowledge of its world and what ontologies can be used to specify the knowledge; the second deals with agent interaction protocols and how these may be formalised. These are set in the discussion of a supply-chain formation.

Like many of the chapters in this volume, **Chapter IV** forms almost a mini-book (at 50+ pages), but Gluck and Fulcher give an extensive review of automatic speech recognition systems covering pre-processing, feature extraction, and pattern matching. The

authors give an excellent review of the main techniques currently used including hidden Markov models, linear predictive coding, dynamic time warping, and artificial neural networks with the authors' familiarity with the nuts-and-bolts of the techniques being evident in the detail with which they discuss each technique. For example, the artificial neural network section discusses not only the standard back propagation algorithm and self-organizing maps, but also recurrent neural networks and the related time-delay neural networks. However, the main topic of the chapter is the review of the draw-talk-write approach to literacy which has been ongoing research for almost a decade. Most recent work has seen this technique automated using several of the techniques discussed above. The result is a socially-useful method which is still in development but shows a great deal of potential.

Petersson, Fletcher, Barnes, and Zelinsky turn our attention to their Smart Cars project in **Chapter V**. This deals with the intricacies of Driver Assistance Systems, enhancing the driver's ability to drive rather than replacing the driver. Much of their work is with monitoring systems, but they also have strong reasoning systems which, since the work involves keeping the driver in the loop, must be intuitive and explanatory. The system involves a number of different technologies for different parts of the system: Naturally, since this is a real-world application, much of the data acquired is noisy, so statistical methods and probabilistic modelling play a big part in their system, while support vectors are used for object-classification.

Amanda and Noel Sharkey take a more technique-driven approach in **Chapter VI** when they investigate the application of swarm techniques to collective robotics. Many of the issues such as communication which arise in swarm intelligence mirror those of multi-agent systems, but one of the defining attributes of swarms is that the individual components should be extremely simple, a constraint which does not appear in multi-agent systems. The Sharkeys enumerate the main components of such a system as being composed of a group of simple agents which are autonomous, can communicate only locally, and are biologically inspired. Each of these properties is discussed in some detail in Chapter VI. Sometimes these techniques are combined with artificial neural networks to control the individual agents or genetic algorithms, for example, for developing control systems. The application to robotics gives a fascinating case-study.

In C**hapter VII**, the topic of structural health management (SHM) is introduced. This "is a new approach to monitoring and maintaining the integrity and performance of structures as they age and/or sustain damage", and Prokopenko and his co-authors are particularly interested in applying this to aerospace systems in which there are inherent difficulties, in that they are operating under extreme conditions. A multi-agent system is created to handle the various sub-tasks necessary in such a system, which is created using an interaction between top-down dissection of the tasks to be done with a bottom-up set of solutions for specific tasks. Interestingly, they consider that most of the bottom-up development should be based on self-organising principles, which means that the top-down dissection has to be very precise. Since they have a multi-agent system, communication between the agents is a priority: They create a system whereby only neighbours can communicate with one another, believing that this gives robustness to the whole system in that there are then multiple channels of communication. Their discussion of chaotic regimes and self-repair systems provides a fascinating insight into the type of system which NASA is currently investigating. This chapter places self-referentiability as a central factor in evolving multi-agent systems.

In **Chapter VIII**, Beale and Pryke make an elegant case for using computer algorithms for the tasks for which they are best suited, while retaining human input into any investigation for the tasks for which the human is best suited. In an exploratory data investigation, for example, it may one day be interesting to identify clusters in a data set, another day it may be more interesting to identify outliers, while a third day may see the item of interest shift to the manifold in which the data lies. These aspects are specific to an individual's interests and will change in time; therefore, they develop a mechanism by which the human user can determine the criterion of interest for a specific data set so that the algorithm can optimise the view of the data given to the human, taking into account this criterion. They discuss trading accuracy for understanding in that, if presenting 80% of a solution makes it more accessible to human understanding than a possible 100% solution, it may be preferable to take the 80% solution. A combination of evolutionary algorithms and a type of spring model are used to generate interesting views.

Chapter IX sees an investigation by Verma and Panchal into the use of neural networks for digital mammography. The whole process is discussed here from collection of data, early detection of suspicious areas, area extraction, feature extraction and selection, and finally the classification of patterns into 'benign' or 'malignant'. An extensive review of the literature is given, followed by a case study on some benchmark data sets. Finally the authors make a plea for more use of standard data sets, something that will meet with heartfelt agreement from other researchers who have tried to compare different methods which one finds in the literature.

In **Chapter X**, Khosla, Kumar, and Aggarwal report on the application of particle swarm optimisation and the Taguchi method to the derivation of optimal fuzzy models from the available data. The authors emphasize the importance of selecting appropriate PSO strategies and parameters for such tasks, as these impact significantly on performance. Their approach is validated by way of data from a rapid Ni-Cd battery charger.

As we see, the chapters in this volume represent a wide spectrum of work, and each is self-contained. Therefore, the reader can dip into this book in any order he/she wishes. There are also extensive references within each chapter which an interested reader may wish to pursue, so this book can be used as a central resource from which major avenues of research may be approached.

Professor Colin Fyfe
The University of Paisley, Scotland
December, 2005