Designing Intelligent Systems

An Introduction

Designing Intelligent Systems

An Introduction

Igor Aleksander

Kogan Page

First published 1984 by Kogan Page Ltd
120 Pentonville Road, London N1 9JN

British Library Cataloguing in Publication Data

Aleksander, Igor
Designing intelligent systems.
1. Artificial intelligence
I. Title
001.53'5 Q335

ISBN 0-85038-860-0

Printed and bound in Great Britain by
Billing & Sons Limited, Worcester

Contents

Preface

This book is based on a series of lectures delivered as part of a multi-disciplinary course to a class of first-year undergraduate students at Brunel University, England. The students were attending courses in electrical engineering, psychology and information systems management.

The nature of the course was challenging because it introduced a novel and ill-defined subject to students with virtually no knowledge of the underpinning topics: mathematics and computing. Consequently, the book is aimed at a very broad readership and truly represents an introduction to this new and absorbing subject. But perhaps the book also aims to achieve a little more than this. Being basic in approach, it has been necessary to include in it only those concepts that are truly fundamental. As the subject in itself is still very much in its infancy, this has meant concentrating on issues that are largely independent of technological 'nuts and bolts'. After much sifting, these fundamentals appeared to centre on systems that help to clarify vague ideas and so allow us to begin to understand the sense in which a system might be said to be 'intelligent'.

Modelling techniques were the next fundamental issue to be considered, and here aspects of automata theory and mathematics have been chosen as foundation topics. The nature of computer programs is clearly central to the design of intelligent systems, so that some well-known programming schemes in this area have been included.

Two chapters on current applications have been included so as not to leave the reader with the impression that the subject is purely theoretical. Finally, there is a glimpse into a possible future, and a description of a possible algorithm that would 'get to know' its user.

Introduction

Intelligent systems as a cornerstone of information technology
We live in an era in which extraordinary claims are being made about the achievements of computers. Every few days the media report on some new super-intelligent feat that someone in an artificial intelligence laboratory somewhere has made. Yet another Grand Master at chess might have been defeated by a machine, or a computer might, without human intervention, have composed a new symphony that sounds just like Beethoven. One even hears of programs that are likely to do away with the human general practitioner.

Unfortunately, when examined closely, such systems only create an *illusion* of intelligence. Computers work very fast, and it is this speed that lies at the heart of the illusion. The all-powerful chess-playing machines of today are really no better than that built by Johann Nepomuk Maelzel in 1830. This machine thrilled audiences by beating its human challengers at chess until it was exposed (by Edgar Allan Poe, among others) as a fake. There was a man hidden inside it! The man hidden inside the modern chess-playing machine is the programmer. He has simply stored some good chess-playing principles, and offset the lack of flexibility of such a set by finding fast ways for the machine to look a long way ahead.

Fortunately, these illusions are being recognized as such and a new science of intelligent systems design is beginning to take shape. The ultimate aim of this science is to create machines in which intelligence is measured not by the way it outsmarts human users, but by the excellence with which it meets human needs. If there is any value to what is known as *advanced information technology* (or *informatics*), it will be to spread the benefits of computing power to ever widening populations.

Breaking away from programming: the task for intelligent systems

The ultimate line of attack is the machine that makes no demands of its user so far as programming is concerned. Clearly, such a machine will have to see, hear and understand. To do this it will need to spend some time in negotiating the meaning of the every-day things in life, so as to provide a bedrock of knowledge on which to build. This will have to be an on-going process, as the machine and its user must continually be updating their knowledge of each other. This book is about the sort of design principles that need to be understood if such systems are ever to be brought to fruition.

Artificial intelligence is a branch of computer science which clearly demonstrates that computers are capable of tackling logical as well as numerical problems. It is not the intention of this book to provide a manual of artificial intelligence, but parts of the book do introduce artificial intelligence techniques and expert systems. The latter are examples of practical applications of artificial intelligence which enable a computer-user to interact with a machine that contains a large collection of facts and rules. The perspective of this book is such that it considers artificial intelligence as being just one component among many that are being used in the quest to provide the machine that learns about its user, and that makes its computing power available without demanding that the user should learn a set of programming techniques.

An overview of the book

Clearly, no intelligent system of the kind outlined above exists in practice. Therefore, it is not possible to take a descriptive approach in discussing systems design. Rather, this book is directed at resolving an ideal intelligent system into its potential elements. It is believed that these elements now exist in computer science, electronics and mathematics.

AN INTRODUCTION TO SYSTEMS

Close examination has been made of the meaning of words, so that the potential for misunderstanding, even when dealing with key terms such as *intelligent* and *system*, is reduced. This definitional attitude is the subject of Chapter 1. After a brief look at previous attempts at defining *systems theory* as a body of knowledge that provides unification, the opposite view is taken: that systems design must involve an appreciation of a variety of non-overlapping methods and theories that have a potential for contributing to the target design. A system is seen as a way of isolating an element of the world so as to be able to study it without interference from other elements. The chapter introduces a systematic way of measuring the attributes of systems by a framework of constructs (such as *manufactured* and *natural*).

The chapter ends with a discussion of a novel construct: *intentionality* and the lack thereof. This is very much a human construct and deals with our ability to relate to other people and objects, by understanding inwardly their likely behaviour. This is thought to be the key construct that will distinguish between illusory and real intelligent systems.

AUTOMATA THEORY

The subject for Chapter 2 is quite specific in that it introduces automata theory, the modelling technique used for dealing with systems having inner complexity. There is no doubt that intelligent systems have considerable inner complexity, which this approach helps to simplify. The power of automata (used in the formal sense) is that they provide a graphical way of visualizing this inner complexity and, indeed, of being able to construct diagrams to illustrate a large variety of inner modes of behaviour. A further pointer to the fundamental nature of formal automata is that there exist totally mechanical ways of turning the graphical representation (known as a *state structure*) into either hardware or software. This means that a system may be designed in the first instance without having to take into consideration such details as the language in which the programs will be written, or the type of chips that will be used in the hardware. After considering a design example at this abstract level, the modelling power of automata is broached again by means of two very different examples. The first is couched in terms of a game of cards played by three people using three strategies. It illustrates the potential of the model not necessarily for card games only, but also for similar situations that could occur in business management, stock markets and strategic studies. The second example is related to psychological modelling where an automaton model is used to provide an explanation of dreaming. This is typical of a class of modelling problems, highly pertinent to intelligent systems, for which only vague theories exist within psychology.

MATHEMATICAL MODELLING

The mathematical basis for informational systems is quite different from that of more classical branches of engineering and physics. In a curious way, the division neatly coincides with another division in mathematics itself: classical and modern. Classical mathematics more or less ended at the turn of the century, although the rumblings of modern theories go back more than two hundred years. It is not so much the subject matter that has changed between the old and the new, as the way in which the *use* of mathematics has been reconstrued. Chapter 3 starts with a review of these changes and shows that modern mathematics is more about logical deductions from a set of *abstract* facts than about predictions of events in the real world. Although this sounds rather devoid of purpose it bears a strong affinity with the operations carried out by any informational machine, whether it be an abstract automaton or a computer program. The logical steps that are found in mathematics between axioms and theorems are analogous to the logical steps taken by an informational system during the course of its operation. In Chapter 3 there are examples of the properties that are looked for in modern mathematical systems, and the way in which these may affect information system models is discussed. In particular, the semi-group is considered: this is the mathematical twin of the finite state automaton discussed in Chapter 2.

ARTIFICIAL INTELLIGENCE

At this point the book turns its attention to the word *intelligent.* Probably the only technical use made of that word previously was in artificial intelligence. This is introduced in Chapter 4 which draws attention to an alternative form of modelling to the automata described in Chapter 2. The word 'algorithm' is a synonym to 'recipe', so the algorithmic models described here are recipes that achieve some stated aim. The aim in artificial intelligence is to perform tasks on a computer which, if done by humans, would be said to require intelligence. Typical are programs that play games, solve problems, and begin to get to grips with the extraction of meaning from simple sentences input into a computer in natural language. Scientists researching artificial intelligence have, as far as possible, tried to make their work independent of the details of the machine on which their programs run or of the particular computer language they use. In this sense the algorithms presented become automata of specific kinds, and departures from classical automata add interest to this work.

USE OF INTELLIGENT SYSTEMS

In the 1960s the major expectations for artificial intelligence were in the fields of automation and space exploration. Making robots more intelligent was the object of much work in the field of automation, while for the space programme the creation of exploration machines with on-board intelligence was the standard objective. In fact, neither of these aims have yet been achieved, the latter for political reasons in the sense that expenditure on space exploration was drastically reduced in the USA in the 1970s, while the former faltered for technical reasons. It is still likely that one of the major outlets for pragmatic intelligent systems will be the factory floor. In Chapter 5 some of the gaps between the needs required for automatic production and the solutions offered by artificial intelligence come under scrutiny. Central to this argument is the impact of the industrial robot. While intelligent information processing has grown away from industrial problems, the mechanical aspects of robots have developed along much more practical lines. For instance, mechanical manipulator arms, driven by simple repetitive programs lodged in largely non-intelligent computers, are now commonplace.

EXPERT SYSTEMS

Probably the most successful application of artificial intelligence in the commercial sense is the expert system. This is almost a direct spin-off of classical work in artificial intelligence and cashes in on techniques for making logical deductions by means of a program. The logical relationships pertaining to a particular area of expertise (such as the diagnosis of disease from a set of symptoms) are fed into the machine by an expert. Such information can subsequently be used by non-experts who interrogate the machine in near-natural language. This technique is examined in Chapter 6 and some of the more successful applications are presented as case studies. Attention is also drawn to the limitations of expert system methods, as this provides important pointers to possible future research trends.

A LOOK FORWARD

It seems important that attempts are made to break away from some of the limitations of artificial intelligence. In Chapter 7 the accent is shifted from the coldly logical action of a computer model of stored knowledge to a situation where programs and systems begin to develop a better understanding of their users. This simultaneously affords to the user a better understanding of the machine with which he is interacting. After all, the ultimate mark of a truly intelligent machine may be in its ability to understand. This implies the use of methods that can cope with the fuzziness of the real world and are able to create idealized representations of such fuzzy data. This field is currently being researched and contains many unsolved problems. An introduction is made to one particular approach in such research and this is related to the possible changes that may be expected in artificial intelligence in the future.

This book then is directed towards those who are embarking on courses of study or self-education in the emerging field of information technology. No prior knowledge of computers, engineering or mathematics is required. An agile and open mind will suffice. Looking towards the next century, those who are new to the field now will experience the most enormous changes in engineering techniques, and these will both impinge on and guide the evolution of intelligent systems. The central aim of the book is to try to anticipate those design principles that might provide a framework for this rapidly advancing process of technical invention.

CHAPTER 1

What is a system and how can it be intelligent?

Historical notions

It may seem excessively pedantic to question the meaning of simple words, when it is clear that the target for discussion is specific. The trouble is that use of the word *system* evokes a wide variety of notions depending on who is using it and who is hearing it. One of the central points of this book is that the design of intelligent systems is a multi-disciplinary process. The designer of an intelligent system will have to keep a very clear head about his procedures. He will need to take notice of systems in industry, as well as those that occur in nature. Further, he will have to understand the interfaces between human and mechanical systems.

Engineers and scientists are very fond of using the word system. They often extol the virtues of the 'systems approach'. What is generally meant is that a methodical approach is being used. It has also been competently argued that there are surprising similarities between ways in which totally different systems operate. The eminent biologist Ludwig von Bertalanffy (1973) probably coined the phrase *general system theory* and published an important book on the subject. The core of this theory is that there exists a body of tools (mainly mathematical) that may be applied in a variety of seemingly disconnected fields. For example, topography (the mathematical way of specifying the way in which things are connected) finds application both in the design of transportation systems and the design of electronic networks.

Although this book does not shy away from looking for such similarities, it tries not to lose sight of the alternative point of view, that the *differences* between different kinds of systems are vitally important too. If one accepts that at least one role for an intelligent system will be to provide a better match between man and machine, its designer must be fully aware of the differences between these two systems.

Systems: a way of partitioning complex worlds
It is often a good thing if the technologist can choose his words so as to have a meaning close to its colloquial meaning. What then does the dictionary make of the word 'system'? There are several definitions, ranging from the tangible to the abstract.

At the most tangible level there is:

> 'Anything made of parts placed together or adjusted into a regular, connected whole.';

while at the most abstract level:

> 'A full and connected view of some department of knowledge.'

Worth noting in these definitions are the words 'together' and 'connected'. The feeling is one of parts interconnecting to achieve some specific purpose. There is an implied cohesiveness among these parts, with a much looser connection to the rest of the world.

A good way of intensifying this understanding is to test it with some specific examples. Take a motor car engine, for example, this has all the markings of a system. Even though a connection is required via the drive shaft to the rest of the car, there is a far tighter coupling between its own internal parts. Now take a book; even though it has the appearance of a system, in the sense that the pages are the parts and the binding provides the togetherness, we feel uneasy to dub a book in this way. The feeling of unease stems from the fact that the pages are really not interacting in a functional way, the binding being there for the convenience of merely grouping *parts of a whole* together. So in some sense a book is a system, but we note the need for expressing its functional complexity in some way.

A further complication about a book is that its contents may well be a substantial system under the second dictionary definition. We shall return to this 'inner' and 'outer' character of functional things often in this book. For now, suffice it to say that it is often the case in systems that 'use' information, that the outer part forms a physical system, whilst the inner part uses an informational one. These two systems are totally different in kind, and yet are intimately coupled to one another. A typical example of such a system is the digital computer; this has circuits intimately interconnected to form a system (the hardware), while the associated software program is a totally different informational system of interconnected instructions and computer-language statements. And yet, to achieve the purpose of the program it is essential that these systems work in perfect unison.

Take now what is probably one of the most complex of systems: the human being. Not only does the inner/outer connection hold (philosophers call it the *body/mind problem*), but there are also several other systems required to keep the whole human functioning. These are chemical (eg digestion), hydraulic (eg blood flow) and, indeed, there are many others on which the reader may wish to speculate.

All this suggests an important point about systems. They help to determine ways in which the things of the world are categorized. The point is that not only does this division occur at the level of physical space, but also at levels of function that require a considerable degree of abstraction to understand them. For example, it may be quite possible for a neurophysiologist to plot out the exact physical connections of cells in the brain, but unless he knows about the history and experience of the particular brain he is considering he cannot predict or even ponder on the likely way in which the human owning that brain is likely, for example, to answer questions about which is his favourite dessert. This is what we mean by abstraction. To predict informational behaviour of the dessert kind, the predictor must know something less observable than cells: the likes and dislikes of the owner of that brain.

There are cases where the inner system may be studied almost independently. This is particularly true of computer programs. A chess-playing program, for example, may be studied as an entity apart from the hardware of the computer on which it will run. Indeed, a very important question that will be raised later in this book is whether, having studied a chess-playing computer program to some depth, the system is valid both in the context of a machine and a human being.

The conclusion reached is that the word 'system', outside of its important role of categorizing things, should not be