The Intelligent Micro

Artificial intelligence for microcomputers

Noel Williams

The Intelligent Micro

Artificial intelligence for microcomputers

Noel Williams

McGRAW-HILL Book Company (UK) Limited

London · New York · St Louis · San Francisco · Auckland · Bogotá
Guatemala · Hamburg · Johannesburg · Lisbon · Madrid · Mexico
Montreal · New Delhi · Panama · Paris · San Juan · São Paulo · Singapore
Sydney · Tokyo · Toronto

Published by
McGRAW-HILL Book Company (UK) Limited
MAIDENHEAD · BERKSHIRE · ENGLAND

British Library Cataloguing in Publication Data

Williams, Noel
The intelligent micro: artificial intelligence
for microcomputers.
1. Artificial intelligence – Data processing
2. Sinclair QL (Computer)
I. Title
001.53'5'0285404 Q336

ISBN 0–07–084770–3

Library of Congress Cataloging in Publication Data

Williams, Noel.
The intelligent micro.
Includes index.
1. Artificial intelligence – Data processing.
2. Microcomputers. 3. Sinclair QL (Computer) –
Programming. 4. Basic (Computer program language)
Q336.W55 1985 001.64 85–12486
ISBN 0–07–084770–3

12345 LT 88765

Printed and bound in Great Britain by Latimer Trend & Company Ltd, Plymouth, Devon

Preface

If you want to learn about a fascinating new application of computers or to begin to learn the principles behind some new computing skills, then this book will do both things for you. It provides a wide and detailed introduction to artificial intelligence (AI) which involves many novel approaches to programming. It can also result in programs which are more satisfying to use and develop than your average game or database because the computer behaves, in some measure, like a human being.

The book is built around 13 programs which use artificial intelligence techniques, plus some additional routines and algorithms. The programs are written in Sinclair SuperBASIC, which is the language used on the QL. This is one of the most advanced BASICs around and is therefore eminently suitable for this kind of work. However, all the programs have been designed to be as easy to understand as possible, and conversion to other BASICs is elementary. In fact I recommend that owners of machines other than the QL carry out such conversion. Two things will be learned: (a) why a highly structured language like SuperBASIC is desirable for AI programming; and (b) that AI routines can be written for just about any machine you care to name. Conversion notes from QL to BBC, C64 and Spectrum are given in Chapter 2.

All programs are tested and working and can easily be adapted to your own purposes. The programs are mainly used to illustrate a wide range of topics which are becoming increasingly important in modern software design, rather than to give you fully operational robots or complex intelligent databases (either of which would require a book much longer than this one). Their main aim is therefore to introduce you to the fascinating and varied fields of work comprising artificial intelligence by showing the range of human capabilities that machines are beginning to possess. Particular attention is given to text processing and natural language because this is an area which everybody can understand. Language is crucial for making software friendly, so if you want educational programs that can chat to the student, databases that can be questioned using ordinary English, or adventure games with commands several sentences long, you will find plenty here to help you along.

But artificial intelligence is a new science and constantly breaking new ground. This means that you will find no easy solutions to real problems which require intelligent software, and in many cases programmers have only just begun to understand the nature of the problems they are exploring. The achievements of AI come nowhere near its potential or its aims. Although there is no such thing as an intelligent machine, this book can show you some of what is involved in creating one. A bibliography is provided so that you can pursue those areas of most interest at your leisure. You should always remember in AI research, however, that the field is so new that your ideas may well be original and it is perfectly possible that you, sitting at home with a micro, will develop AI ideas with capabilities and cleverness in an area not even contemplated by researchers at IBM or MIT.

As no system can yet claim true intelligence, and no book can show you how such a complete system might be achieved, my aim has been to give programs which begin to make micros human-like. Intelligence is a complex and multi-faceted phenomenon and it is only by individual exploration of each of the facets of human intelligence and behaviour that we can gradually advance towards a completely intelligent system which combines all such features (assuming such a thing is desirable). By the time you have finished the book you will be able to write programs that can correctly analyse sentences, write poetry, create jokes, make decisions on complex evidence, recognize threats and make inferences—programs that begin to think, to learn and to 'understand'. But, more important than this, you will also understand how such programs work and why they are limited, so that you can go on to develop better and more capable programs of your own.

It would be a mistake to pretend that a program can be fully intelligent. All we can do at present is simulate various aspects of human intelligence with some degree of accuracy. You will find that even in those areas where AI has had major success, such as with expert systems, researchers can be puzzled by their own successes. We are only just on the fringes of an exciting new form of computing and there are an enormous number of puzzles which need solving. You could be involved in the solution of some of them.

This book would still be lying on my desk were it not for the affection and hard work of Carrol, who deserves all the thanks she can get.

[Note: The listings in this book are taken directly from running programs. However, the printer used has shown each # sign as a £ sign. Therefore, wherever a £ sign is encountered it should be replaced with a # sign.]

Contents

1 What is artificial intelligence?

1.1 Introduction

As computers, microprocessors and microcomputers become more and more important in our lives we hear more and more about artificial intelligence, about intelligent knowledge-based systems (or expert systems), about the Japanese fifth generation of microcomputers and the Alvey report on Britain's need for an equivalent research project. What is all the fuss about? And why is it happening now?

One of the reasons for the rapid growth of interest in artificial intelligence (let us agree to call it AI for the rest of this book so that you do not get fed up with the persistent phrase), is the new set of microcomputers becoming available, such as the QL, Apricot and IBM PC, with faster and larger processors and much larger memories than the 64K available to the tired old 6502 and Z80 micros we are so familiar with. Because micros with 128K, 256K or even half a megabyte of RAM are now coming into the price range of anyone who can afford a second mortgage, interest is growing in applications which can use that massive extra RAM. Partly this is because people with large computers tend to feel that they are not using them properly (whatever that means) if they are not using all the extra memory, and partly because developments in AI have always been held back by the expense of suitable hardware. Now most micro enthusiasts have the potential to begin developing their own original investigations of the limits and rewards of AI programming. Just as an enormous series of developments has taken place in hardware and software in the wake of the video games boom, so it seems likely that a similar 'explosion' of interest will occur when people sit down to think 'what worthwhile things can I do with my micro?'

The main aim of this book is to introduce you to some of the concepts of AI but without being too theoretical and dull. On the one hand you will learn something about the general aims, achievements (and also the major problems) in AI programming; on the other hand you will find a number of illustrative programs that do somewhat clever things, most of which can easily be incorporated into larger programs or fitted together to construct more versatile routines. Although AI can be a difficult area of computer study, that does not mean that it has to be all the time, nor that it cannot be enjoyable. AI is above all things about people and that can make any subject interesting and entertaining. It also enables the micro enthusiast of today to enjoy some of the thrill of pioneering and discovery that fired the first home computer hobbyists five years ago, for AI is a very new kind of computing where new developments occur almost every day.

So the aims of this book are three:

1. to introduce you to a number of aspects of programming which are becoming increasingly important and which are collectively known as AI;
2. to show how BASIC can be used for writing AI routines, illustrated through one of the best BASICs available;
3. to illustrate some simple AI techniques and show how some of them can be adapted for use in your own programs. In particular we will pursue detailed investigation of various aspects of natural language and text processing because these are areas which can greatly enhance serious business programs, educational software and games like adventures and simulations.

One thing this book does not spend much time on is the kind of AI used in abstract games programs like chess and draughts. These routines consist largely of different search strategies. Search is an important concept, not only in AI but in other branches of programming, but the searches employed by most abstract games are often specific to the kind of game and as such are only of interest to writers of such games.

1.2 Defining intelligence

So what exactly is AI? There are many definitions, none of which are really comprehensive because just about anything that a computer does could be regarded as 'intelligent' in some sense. Does not something that can do quadratic equations, create moving graphics, measure the length of words and carry out the instructions of a human being seem intelligent to you? Well, you say, it depends what you mean by 'intelligent'. And that is the crux of the problem of definition. The nature of AI depends on what you think of as the nature of intelligence. Some people think that intelligence is no more than a complex series of rules for decisions and that if we knew enough about those rules and decisions we could write programs to carry them out, and the machines carrying out those programs would be exactly as intelligent as the human beings who originally held those rules and made those decisions. Others argue that 'intelligence' is something uniquely human, perhaps even God-given, and therefore it cannot be discovered, described or defined because it defines us. Then there are a number of people who are somewhere in between these two extremes who think that a lot of the activities we call intelligent behaviour actually are rule-governed activities, just like computer programs, but that maybe some intelligent activities are not like that and perhaps we will never know about some activities. You can even argue that it is the nature of intelligence that it cannot understand itself, so even if human intelligence is no more than a series of rules, human beings will not be able to discover them.

For example, it is reasonable to say that understanding the word 'RUN' requires intelligence. We would know that someone understands this word if they carry out the operation that the word represents. So Sebastian Coe is intelligent because if you say 'RUN' to him he will lope off in a cloud of dust. Similarly your micro will carry out a series of activities if you type 'RUN'. But only if it has a program in it. Sebastian Coe shows his intelligence (though the lazier among us might prefer to think of it as stupidity) by carrying out a series of complex motor activities which, among other things, stop him from falling over and make sure that he gets from point A to point B in as short a time as possible. He is showing that he has a certain 'program', i.e., he has learned to act in certain ways in response to certain instructions. A program in a micro can be seen as the current state of learning of that machine and we can measure its intelligence by the accuracy with which it carries out its instructions, the efficiency with which it achieves the intended goal of those instructions and the number of inadequacies it has or errors it makes on the way.

From this point of view anything which carries out instructions, or responds to a signal, has some kind of intelligence. We can think of a kind of hierarchy of intelligence with stones pretty close to the bottom, plants slightly higher, insects a little higher, mammals rather higher and human beings (you and me anyway) somewhere near the top. On this kind of scale most micros would be between plants and insects. Plants respond to stimuli in a very restricted and highly conditioned way. If the light comes predominantly from the west, all the plants grow to the west. Maybe the light source is a lava flow which will frazzle the plants when they grow big enough, but they are not responsive enough to 'know' that.

Insects, on the other hand, not only respond in rather mechanical ways to stimuli, but they can create stimuli (i.e., they can 'give instructions'), they can learn and they can develop complex patterns of behaviour which have meanings, such as the dances that bees use to indicate a source of pollen. Most micros that I have been introduced to are not as clever as these bees. Micros can be given instructions and can carry them out, but they cannot instruct themselves. They cannot communicate with each other about things of interest to themselves. They cannot make decisions about which actions are in their best interests or in the interests of microcomputerdom in general. And where they are beginning to do such things it is not because they have 'discovered' how to do it but because humans have made the discovery and turned it into language sufficiently simple for the micro to understand.

So intelligence probably does not simply involve responding to stimuli and executing learned 'programs'. It involves other things as well. Psychology is the science which attempts to discover what these other things are and AI is one branch of psychology. Computer scientists are involved, linguists are involved, engineers are involved, philosophers are involved, but the main thrust of AI research comes from psychologists, because psychologists want to discover how human minds work. One way to make this discovery is to make a model of the human mind, just as we might demonstrate how a steam engine works by making a model, and AI is really about making models of human intelligence.

In this case the question 'Is the machine intelligent?' is almost meaningless. If you build a working model of a steam engine which does not look like a steam engine, but works like a steam engine, can do everything that a steam engine does and is a perfect replica in every significant detail of a real steam engine is your model a steam engine or not? Who cares? To all intents and purposes it is exactly like a 'real' steam engine. It behaves like a 'real' steam engine. If in some deep philosophical sense it is not a 'real' steam engine because it is smaller, or because it is not being used like one, or because it was not designed to be one, that hardly matters. In the same way it does not matter if an 'intelligent' dog actually thinks in a way we could not possibly understand which is nothing like our own intelligence. If the dog recognizes his mistress's voice, if he sits when told, if he fetches the paper when it is delivered and pines when his mistress is away, then he has a degree of intelligence. He understands what is going on around him and acts accordingly. The question of whether his intelligence is like our intelligence is unimportant.

So the [illegible] boils down to a debate about the definition of words. When we talk about an intelligent machine, what we mean is a machine which, as far as we can tell and within a limited area, acts as a human being would act in the same situation and with the same information. To understand this, just remember the Turing test. Alan Turing, an eminent mathematician in the thirties and forties, devised a test to determine whether a machine should be called intelligent or not. A human subject is placed in a room with two identical teletype terminals. All communication with the outside world takes place through those two terminals. Communicating with one terminal sits another human being, while the other is controlled by a computer. The human in the room has to determine which teletype is controlled by the computer and which by the human being. If he or she cannot decide, then, in every way that matters, the machine is as intelligent as the human being; it has passed the Turing test.

No machine has ever passed the test. But some have made some people think that they were human beings in special circumstances. For example Joseph Weizenbaum developed a program called 'Eliza' which simulates the responses of what is called a 'non-directive' psychiatrist, i.e., one who just tries to encourage patients to talk without passing any comment on what is being said. When his secretary first tried it out, she requested that he leave the room because she was holding a private conversation. As far as she was concerned, and even though she knew she was 'talking' to a computer program, it was an intelligent process which was 'taking her seriously'.

If we heavily restrict the kind of communication that is permitted through the teletype in applying a Turing-like test; for example, by restricting the topic of conversation, many machines and programs can be regarded as intelligent in a limited sense. What this means is that AI exists in a restricted form. There are machines which are intelligent in highly restricted areas. Psychologists and others have discovered some of the rules of some kinds of intelligence. Machines can imitate certain aspects of human behaviour. But no machine has anything like the flexibility, versatility or range of the dullest human being. However, such machines may not stay in the realms of science fiction for long and you may even be involved in bringing them to life—if you want to.

So there are only really two ways of defining AI which make any sense. One is to list all the things which people have called intelligent activities and which research has shown machines can model, and the other is to list all those things which people regard as intelligent behaviour which we might want machines to model but which no one has yet successfully made them model. The reason the last list is difficult to compile is because no one yet has a satisfactory definition of human intelligence or a satisfactory list of all the things that a human being can do which we might regard as intelligent. If we look at it from the ordinary person's point of view for a minute we might say that an activity appears clever if the person, creature or machine carrying it out satisfies one or more of three criteria:

1. It is doing something which we would not expect it to be able to do.
2. It is doing something we could not do or would find difficult to do.
3. It is doing something apparently with a purpose, intended to achieve a particular goal.

Criterion (1) simply means that our understanding or idea of the machine or creature is inadequate. Dogs might appear clever because we have a very limited idea of what a dog can or should be able to do. Criterion (2) simply means that we are inadequate in some sense—we are slower or weaker or illogical or have incomplete knowledge and so have not the ability to carry out the action. If we had the ability or the knowledge we might not regard it as special activity at all. For example if you understand machine code you regard it as easy, whereas everyone who cannot separate their bits from their bytes thinks you are particularly intelligent. In other words these definitions of 'intelligence' are relative to the person who is using them. When you know the secret you no longer regard it as intelligence.

Thus it is argued that computers that beat chess masters appear intelligent only to those who do not know the chess strategies and the programmed logic involved. You might argue in a similar way that chess masters appear intelligent only because we do not know what it is they are doing when they thrash us poor mortals at the game. But if we did know we might not regard it as intelligence at all. It seems clear that some of the things that a grand master does are the same as the things a computer does when deciding which piece to move—both 'know' something about previous games, both examine a number of different moves and look at their consequences, both try to balance the advantages for themselves against the advantages for the opponent. It is also clear that some of the things they are doing are different—the computer (in principle) might evaluate every possible move resulting from a particular position, whereas the grand master might 'feel' that a particular move looks promising and make his move on the strength of this feeling.

Criterion (3), however, is a useful one. Intelligence involves purpose and the purpose must come from the intelligent creature itself. This is a stumbling block with machines because they cannot really have any intentions. However they can be given goals, they can be given purposes which are extensions of the purpose of the programmer, and to this extent they have intelligence. The whole idea of 'strategy' depends on having a goal—a good strategic decision is one that brings you nearer the goal and a bad one is one that takes you away from it. Thus psychologists and AI researchers are very interested in goal-directed activity, the nature of strategy and 'why people do the things that they do'. Of course a major area for exploring the nature of strategic thinking is in games, and this is one area which can benefit from AI research.

If we now try to list the areas which have been explored with some success as a way of trying to define AI, the following list would cover most things:

1. pattern recognition and perception
2. robotics
3. creativity
4. problem solving
5. decision making
6. learning
7. natural language understanding
8. personality theory
9. human-computer interaction and the man–machine interface.

This is quite an assortment of areas and there is not room to cover them all with the same degree of detail in one book. However, you would probably agree that a machine which could create new ideas, hold a normal conversation, recognize the person talking to it, solve a variety of complex problems and make decisions about possible solutions, and could learn from its mistakes was either intelligent or a very good model of human intelligence. Unfortunately (or perhaps fortunately) there is no such machine and not likely to be one for some time, if ever. Though large claims have been made for work in nearly all these areas, and there certainly have been some notable advances, when you actually examine the achievements of AI they are not very spectacular. But they are interesting and there are several useful and fascinating areas of development. In particular it seems that if we are to have 'user-friendly' machines which most people will be happy with in a wide variety of situations, they need some of the routines provided by such work. So you might want to learn about AI simply to improve the so-called human interface between your software and its user. And this might apply whether you are producing a word processor or a mind-taxing adventure.

2 AI on micros

2.1 Aspects of AI

For the rest of this book we will forget the debate about the nature of machine intelligence and just assume that AI means 'making computers clever'. Whether you think they are actually clever or only models of cleverness is up to you. Let us also assume that you want to know about AI in order to write programs which are (or have some of the characteristics of) real human beings. We shall explore as many aspects of this as possible, as much to discover what the limitations of AI are as what can be done, and to see what needs to be understood in order to go past those limits. The rest of this chapter gives a brief description of each of the aspects of AI just listed, then examines some of the programming concepts involved and how they can be implemented in Sinclair SuperBASIC, and gives some notes on conversion to other BASICs. Then Chapter 3 begins the real analysis of AI, with chapters dealing with the topics described in this chapter—creativity, data storage and coding, natural language understanding and production, personality theory, pattern recognition, semantics, heuristics, learning, decision making and expert systems.

2.2 Robotics

Though robotics involves important work in AI it is not covered by this book as it is really about microcomputer control and sensors. The study of robotics is wide ranging, but broadly speaking the intelligence of robots can be described as very limited. Either a robot can only do exactly what it has been programmed to do, with no deviation or variation, such as the robot workers on assembly lines which exactly mimic human actions they have been programmed to copy, or, where they do have some degree of control over their environment, they are not able to carry out very important or sophisticated actions. In other words, the more useful the robot, the less intelligent it tends to be in practical terms. Understanding robots involves a great deal of knowledge of control systems and much less of AI.

2.3 Pattern recognition

However, in order to carry out any kind of action a robot has to know what it is acting on, and in this respect some robots have quite complex sensory apparatus for recognizing patterns. Some merely carry out actions blindly (literally), that they have been instructed to carry out and will spray the object in front of them in exactly the same way whether it is a car, a QL or a human being. But others can follow infra-red signals, or search for light sources, or recognize quite complex shapes when deciding whether an object is to be manipulated or not. AI has played a large part in describing what human beings pay attention to when they perceive things and in enabling machines to carry out similar operations.

In particular it is important for an intelligent system, whether it carries out control activities or simply provides a friendly interface for a database, to be able to recognize patterns of action in the world around it. The activities we call 'perception' do not simply involve the passive reception of signals, but active interpretation—even at the most fundamental level. The brain is constantly on the look-out for patterns of all kinds (including 'patterns of patterns') and it is this facility which allows it to sort out ambiguous messages and process information efficiently, whether its owner is listening to music in a crowded hall, trying to shoot down a space invader or translating a text from Welsh to Urdu.

2.4 Creativity

Creativity is possessed by all human beings in greater or lesser degree. For example we all have the capacity to produce an infinite number of different original sentences, assuming we would want to do so. We can create new ideas, new objects and new words. For many people artistic creativity in particular is something that machines will never be able to capture. However, as we will see, there are some aspects of creativity at least that can be programmed, and it seems likely that several kinds of relatively structured types of creativity will eventually be available to micros. As in some measure all creative activities seem to involve similar processes, if the basics can be programmed then the QL can begin to write, compose, paint or sculpt—or at the very least to become a useful tool providing ideas for a human being who wants to do such things.

2.5 Problem solving and decision making

These 'thought processes' have been most successfully implemented in so-called expert systems, which we will examine later on. The ability to