Pattern Recognition

Ideas in Practice

Edited by Bruce G. Batchelor

Pattern Recognition

Ideas in Practice

Edited by
Bruce G. Batchelor
University of Southampton, England

Plenum Press • New York and London

Library of Congress Cataloging in Publication Data

Main entry under title:

Pattern recognition.

Includes bibliographies and index.
1. Pattern perception. I. Batchelor, Bruce G.
[DNLM: 1. Biomedical engineering, 2. Pattern recognition. QT34 P315]
Q327.P37 001.5'34 77-12488
ISBN 0-306-31020-1

A Division of Plenum Publishing Corporation
227 West 17th Street, New York, N.Y. 10011

Printed in the United States of America

Contributors

W. A. Ainsworth Department of Communication, University of Keele, Keele, Staffordshire, England

I. Aleksander Department of Electronic and Electrical Engineering, Brunel University, Uxbridge, Middlesex, England

L. A. Amos M.R.C. Laboratory of Molecular Biology, Hills Road, Cambridge, England

B. G. Batchelor Department of Electronics, University of Southampton, Southampton, England

Donald A. Bell Computer Science Division, National Physical Laboratory, Teddington, Middlesex, England

C. D. Binnie Instituut fur Epilepsiebestrijding, Heemestede, Netherlands. Formerly at Department of Clinical Neurophysiology, St. Bartholomew's Hospital, 38 Little Britain, London, England

C. M. Brown University of Rochester, Rochester, New York 14627; Formerly at Department of Artificial Intelligence, University of Edinburgh, Forest Hill, Edinburgh, Scotland

M. J. B. Duff Department of Physics and Astronomy, University College London, England

A. S. J. Farrow Medical Research Council, Clinical and Population Cytogenetics Unit, Western General Hospital, Crewe Road, Edinburgh, Scotland

D. K. Green Medical Research Council, Clinical and Population Cytogenetics Unit, Western General Hospital, Crewe Road, Edinburgh, Scotland

P. D. Green Department of Computing, North Staffordshire Polytechnic, Blackheath Lane, Stafford, England

J. A. G. Halé Mullard Research Laboratories, Cross Oak Lane, Redhill, Surrey, England

D. C. Mason Medical Research Council, Clinical and Population Cytogenetics Unit, Western General Hospital, Crewe Road, Edinburgh, Scotland

E. A. Newman Computer Science Division, National Physical Laboratory, Teddington, Middlesex, England

R. J. Popplestone Department of Artificial Intelligence, University of Edinburgh, Forest Hill, Edinburgh, Scotland

J. R. Parks Computer System and Electronics Division, Department of Industry, Dean Bradley House, Horseferry Road, Westminster, London, England

D. Rutovitz Medical Research Council, Clinical and Population Cytogenetics Unit, Western General Hospital, Crewe Road, Edinburgh, Scotland

P. Saraga Mullard Research Laboratories, Cross Oak Lane, Redhill, Surrey, England

G. F. Smith Neuropsychology Unit, Department of Applied Psychology, University of Aston, Birmingham, England. Formerly at Department of Electronics, University of Southampton, Southampton, England

D. W. Thomas Department of Electronics, University of Southampton, Southampton, England

J. R. Ullmann Department of Applied Mathematics and Computing Science, University of Sheffield, Sheffield, England

Preface

Pattern recognition is a child of modern technology; electronics and computers in particular have inspired research and made it possible to develop the subject in a way which would have been impossible otherwise. It is a rapidly growing research field which began to flourish in the 1960s and which is beginning to produce commercial devices. Significant developments have been made, both in the theory and practical engineering of the subject, but there is evidence of a schism developing between these two approaches. Practical machines have usually been designed on an *ad hoc* basis, with little use being made of advanced theory. It is difficult to provide a rigorous mathematical treatment of many problems pertinent to a practical situation. This is due, in part at least, to a conceptual rift between theory and practice. The mathematics of *optimal* systems is well developed, whereas pragmatists are more concerned with vaguer ideas of *reasonable* and *sufficient*. In some situations, the quest for optimality can constrain research and retard practical progress. This can occur, for example, if too narrow a view is taken of "optimal": the accuracy of a system may be optimal whereas its speed, cost, or physical size may be grossly suboptimal.

The objective of this book is to present a glimpse of the pragmatic approach to pattern recognition; there already exist a number of excellent texts describing theoretical developments. Ultimately, pattern recognition will be judged by its contribution to human society and there can be little confidence in its future unless we continually satisfy our paymasters that the research is bearing fruit. This book is offered as an "interim report," showing the breadth and direction of current research. The full potential of pattern recognition has yet to be realized, but the present indications are that very substantial benefits are to be expected – as we hope to show.

It would be nice to begin by formally defining our subject, but so far no pattern recognition worker has been able to do this succinctly and to the satisfaction of the majority of his colleagues. Pattern recognition is inherently a science of ill-defined concepts. It is part of the broader field of information processing and is devoted primarily to (selectively) reducing information content. Information selection tasks occur, for example, in medical diagnosis, biological classification, television bandwidth compression, picture processing and enhancement, automatic inspection, optical character recognition, speech recognition, and a host of other applications.

The physical origin of the information is unimportant to the theoretician but not to the engineer concerned with converting the data from one medium to another. The physical properties of the source of the data may be well understood and this knowledge may provide important clues as to how the information may best be coded. For example, a motor-vehicle sound may be expected to contain components at a frequency equal to the crankshaft rotation rate. The question of what we measure is often dictated by economic considerations; there is then little or no opportunity to *optimize* the measurements. In medicine, for example, pathology laboratories have invested large amounts of money in auto-analysis equipment, which would not be discarded or modified significantly merely at the whim of a pattern recognition worker (who is often regarded with suspicion, as an interloper, by his medical colleagues). In some applications, the primary information rate is very high: 10^8 bits/s in satellite photography is perfectly feasible. A special-purpose machine is then essential and this may severely constrain the algorithms which can be employed. Algorithms must often give way to heuristics or unproved procedures whose mathematical analyses have not yet been completed. There may be many reasons for this, the most common being:

1. Certain applications-oriented studies generate the need for new techniques which must be developed quickly (by experimentation), without the luxury of detailed mathematical analysis.
2. Factors other than system performance may be critical. These may place severe demands upon the complexity of equipment and the operations which it can perform. The size, weight, speed, power consumption, or cost of a machine may be just as important as its accuracy.

We hope to convince the reader that pattern recognition is primarily an engineering discipline, drawing upon a broad spectrum of skills. Indeed, one of the major attractions for many of its practitioners is that it is a truly interdisciplinary subject, calling upon such diverse areas as electronics, optics, statistics, psychology/ergonomics, programming, and systems engineering. Pattern recognition research is often conducted under a different guise, or in laboratories specializing in some applications area, e.g., medicine, factory automation. This is evident here from the list of contributors. Many research workers "drift" into pattern recognition from other areas, seeing it as a source of new techniques for application to their particular problems. Many of my coauthors feel, as I do, that the pattern recognition community has done little to court these researchers and entice them further into our subject. This book is written with the express intention of doing this. However, we believe that we also have a point to make to our theoretically inclined colleagues specializing in pattern recognition; there are many procedures whose genesis is attributable to pragmatism, intuition, and applications-oriented research, and which require formalization by mathematical analysis. We hope to stimulate our colleagues into devoting their attentions to analyzing "practical" procedures.

In view of the interdisciplinary nature of our subject, we have deliberately tried to keep our discourse at a level which is comprehensible to a new graduate in science of engineering.

Numerous "unseen" people contribute to a book such as this. Many people assisted in the formulation of ideas or encouraged our endeavors with their interest. To my mentors, may I say "Thank you." I should like to express my gratitude to my coauthors who have tolerated my whims and fancies with remarkable patience. It is a pleasure to acknowledge the willing help of my secretary, Miss Julia Lawrence. Finally, I should like to thank my wife, Eleanor, for her unending encouragement and interest.

B. G. B.

Southampton

Contents

METHODS AND MACHINES

1

Setting the Scene

B. G. Batchelor

The comments here are my own and do not coincide in all respects with those of my colleagues. I have tried to present a point of view that is close to the consensus among the authorship, but I realize that it is impossible to represent another person's opinion with exactly the same emphasis and nuance.

1.1. SETS

Some sets of objects can be adequately represented by a few *archetypes* (variously called *paradigms*, *exemplars*, *representatives*, or *typical members*), while others can only be specified by listing every member. Some sets can be defined by a suitable description, but other sets cannot. These two sentences restate the same basic truth in different words. The significance of this will be seen throughout this book, since *pattern recognition* (PR) is concerned with sets.

Consider the sets of animals that children call "cows" (male, female, and eunuchoid). It is difficult to define the limits of such a group, although its members all have certain essential (or nearly essential) properties. For example, cows are, for a child, animals with four legs, a tail, possibly horns. They also eat grass and go "moo." While this is not a scientific definition of domestic cattle, it is sufficient for many purposes, including ours. Despite the inadequacies of this or any other definition, I can recognize cows with great ease and accuracy. By going on many car rides, my children have learned to recognize cows just as well as I can. The only effort needed of me was to show my children enough

B. G. Batchelor • Department of Electronics, University of Southampton, Southampton, England

cows, each time muttering the appropriate word, "cow." My children did the rest. Suppose we could do this with a machine: teach it to recognize not cows but faulty components for a car engine. Such a machine would only need to be shown enough samples of the appropriate part, and it would then learn to recognize "good" (i.e., not faulty) components without any further action from me. (Note that I did not say "build a machine to recognize good components.") This is one aspect of PR. Another important facet is *classification* (or *discrimination*). Even without specifying the limits of the four classes – cows, pigs, horses, sheep – we can, with no special training, classify (almost all) the four-legged animals (in the fields of England). This is the basis of that subdivision of PR known as *pattern classification*. Another topic is called *clustering*. To illustrate this, let us consider subsets of cows. There are many distinct *breeds*, for example, Jersey, Hereford, Ayrshire, which can be identified by the observant person, even without *training*. Of course, a person would not be able to place the correct "label" on a subset (breed) of cows without training, but it is certainly feasible for the untrained observer to identify the characteristics of a group. Could a machine learn to do so? Each type or breed of cow has a distinct archetype. The set "cows" has many archetypes, whereas "Ayrshires" has only one (per sex). Note that the word "learn" has, so far, been used twice. Strictly speaking it should also have been used to emphasize the learning aspects of classification. *Learning is central to all aspects of pattern recognition.*

Let us turn our attention now to a different kind of set, and in particular the set of PR problems. This set cannot be specified as easily as "cows." At the beginning of this chapter two types of sets were mentioned. "Cows" very definitely lies in the first type and can be represented adequately by few archetypes. I am tempted to place "PR problems" in the second type. However, to do so would be to admit defeat; as the editor of a book on PR, I should be able to indicate the common areas of PR problems. This is extremely difficult. Neither I nor any of my colleagues have found a succinct definition of PR which does justice to our personal understanding. PR means different things to each of us. That is why this book contains such diverse contributions, but I hope that the reader will find far fewer than 15 archetypes. (My own feeling is that there are three major forms of *pattern description*, each supporting a group of decision-making techniques. See Section 1.3.3.)

1.2. MOTIVATION

Pattern recognition problems are too diffuse in nature to allow us to define the problem in formal terms. There is, however a deep underlying unity of purpose in PR research: to build machines that will perform those tasks that we have previously regarded as essentially "human." A few examples are given

Table 1.1. Illustrating the Wide Diversity of Pattern Recognition Applications and Input Media[a]

Problem	Application	Medium of input
Reading alphanumeric, typed/handwritten characters (optical character recognition)	Reading invoices, checks, etc.	Visual
Recognizing spoken words (speech recognition)	Speech input to computers Speech control of machines Phonetic typewriter	Acoustic
Distinguishing patients with biochemical imbalance from those who are "normal"	Medical diagnosis	Chemical
Inspecting machined components	Automatic assembly/ inspection	Tactile/visual
Detecting possible people "at risk" for certain diseases	Sociology/psychology/ epidemiology	Questionnaires
Analyzing neurophysiological signals	Electroencephalography, electrocardiology, electromyography, etc.	Electrical

[a]Many more examples are described later in this book.

in Table 1.1, and many more will be encountered in later chapters. Many lay people, when hearing the term PR, automatically turn their attention to visual images. This is far too restrictive, as Table 1.1 indicates. The basic event being observed, the pattern may occur in a visual, acoustic, electrical, chemical, tactile, or more abstract "world." PR has often been regarded as a specialized branch of statistics, which we all accept as a legitimate field of study without specifying where the input data arise. Once the transducing problem has been overcome, there is often a strong similarity in the nature of PR problems arising from any of the "worlds" exemplified in Table 1.1. This is not universally true, and there are two major subdivisions of PR techniques as a result. However, the distinction between these is ill-defined and some problems have characteristics of both, while others do not fit conveniently into either. In fact, when we attempt to apply PR *techniques* recursively to PR *problems*, we fail to gain any clearer insight into the basic nature of the subject: *meta-pattern recognition* is still embryonic. We must resort then to human values and human judgment.

PR is motivated partly by our desire to free human beings from tedious, boring, repetitive, and dangerous work. Why should a person become a "vegetable" through working in a factory, or risk contracting cancer through inspecting or repairing the inside of an operational nuclear reactor? Copytyping is a boring job, especially when the copytypist does not fully understand the text he/she is duplicating. In our attempts to control certain diseases, we use mass-screening techniques, which place boring, mundane "inspection" tasks upon technicians or doctors. There are strong economic motivations for PR as well. It is too expensive to inspect every component of a (mass-produced) motor car manually, but the result is a product which is annoying to the customer and costly to repair under warranty. Copytyping is expensive as well as tedious. The same is true of computer-card punching. Finally, there are certain tasks that people cannot do well under any circumstances. If a person is presented with a table of numbers listing, say, the results of a 10-measurement biochemical assay on each of 10,000 patients, little more than bewilderment will occur. There is little chance that he/she would be able to find a basis for discriminating between "normal" people and those with an "abnormal" profile of measurements. We can list the three major incentives for PR research as:

1. *Social* (i.e., eliminating boring, repetitive, or dangerous work)
2. *Economic*
3. *Improving performance*

1.3. COMPONENT PROBLEMS OF PATTERN RECOGNITION

The factors noted above as motivating PR are at the heart of the subject. They provide the most effective "binding" for our discussion. This must be realized before we begin our discourse in earnest, because the variety of techniques that we shall meet may produce a feeling that the subject is a hodgepodge. However, there are certain identifiable areas of PR which we shall discuss now.

1.3.1. Transducing

First of all, the pattern must be translated into a form suitable for manipulation in a computer. Electronic technology, especially digital computers, provides the processing flexibility and speed that we shall need. At the moment, the only serious contender to electronics is optics, but this is still in a rudimentary state of development, compared to the extraordinary variety of analogue and digital integrated circuits that are now available. Once in electronic form, the data may still be in an inconvenient format. For example, unary- to binary-coded-decimal conversion may be required before the data can be used by our computer. This type of format conversion is a mere technicality and will therefore be dismissed here.