Formal methods of diagnosis in engineering and medicine

By A. J. Lerner*

In this article some problems of formal diagnostics are formulated. Difficulties arising are discussed and a computer approach to medical and technical diagnostics is offered.

The approach consists of organisation of a process of automatic production of a decision rule on the basis of information stored in the set of parameters of correct classification of situations. A list of problems in medical and technical diagnostics is presented; these were solved by the author and his collaborators by an original method of learning, the so-called method of 'Generalised Portrait'.

This article is based on a talk given to the Medical Specialist Group of the BCS in London in April 1967.

1. Formulation of the problem

We use the term *diagnosis* for the procedure of allocating complex situations to different categories and subcategories on the basis of their 'symptoms' taken collectively. If the symptoms used for diagnosis are the same as those used to define the categories, then the diagnostic procedure consists merely in determining the presence or absence of the symptoms, and the problem becomes trivial. The situation is far more complex if the symptoms used for setting up the classification cannot be used for diagnosis, perhaps because they are not fully established or are not observable at the appropriate time, so that any diagnosis must be made on the basis of other symptoms. Problems of this kind often appear in engineering, medicine, economics, sociology and many other areas of contemporary life.

Until recently, problems of this kind were either solved intuitively on the basis of expert experience, or were not solved at all. However, even when expert opinion allocated a particular situation to a certain category the reliability of this kind of diagnosis often left much to be desired. The number of false diagnoses used to be—perhaps still is—too great, with serious consequences; people die, engineering systems fail, resources and efforts are wasted.

Recently, however, formal diagnostic schemes have begun to be evolved and these promise to bring more effective solutions to the problem of classifying complex situations than has hitherto been possible. In many cases these formal procedures enable us to improve the reliability of our diagnoses to a marked extent and also to diagnose situations to which intuitive diagnosis could not be applied. In engineering the frequency of such problems is quite high. They appear in relation to the classification of phenomena in an engineering project—its pattern of performance, its satisfactory or defective operation, its design quality, etc. Engineering diagnosis includes, for example, checking the performance of an engine by its sound, checking a computer by test procedures, establishing the presence of mineral ores on the basis of geophysical and geological readings. The task of medical diagnosis consists in assessing the state of an animal or human organism on the basis of recalled and clinical data. The object is to establish the presence of a disease, the type of disease present, its stage, pathological changes relative to the normal situation, pregnancy, foetal sex, etc.

The problem of classification turns up not only in engineering and medicine, but also in economic, sociology and criminal investigation. For example, if one needs to determine criteria of efficiency for a particular economic system, it is very important to separate more desirable situations from less desirable ones; in production planning, it is important to have a clear picture of the market situation. In sociology, also, there are complex problems of classification involved in grouping members of a particular society into economic categories, grouping members of a population according to their levels of intellectual attainment, or in identifying individuals by aspects of their handwriting or personal appearance. These and similar problems arising in the classification of complex situations by means of large sets of symptoms, though difficult, are extremely important, and the use of formal methods for pursuing them may be both useful and effective. By 'formal' methods we imply those essentially suited for computer processing.

The problem of formal diagnosis, then, consists in setting up formal procedures which are available not only to man but also to the computer, and which make it possible to take decisions with a known degree of certainty allocating a particular situation to one or another predetermined category on the basis of data relating to the whole symptom complex.

2. Difficulties encountered in the diagnostic process

The difficulty of a classification problem depends mainly on the number of symptoms employed and on the nature of the interdependence between them. The difficulty may be very different for man and machine. For example, the classification of numbers into odd and even can be done with equal ease by man or machine since the distinction is based on a single readily perceived symptom; the distinction between men and women based on external appearance is comparatively simple for man, using a few informal symptoms, but presents the machine with serious difficulties. On the other hand man finds it difficult to establish the presence of mineral ores on the basis of geophysical and geological data, whereas this task is solved simply and quickly by the computer.

* *Institute of Automation and Telemechanics, Moscow, USSR*

If we ignore trivial classification problems in which only a small number of symptoms are used in allocating a situation to a particular category, and concentrate on the more difficult problems, which are naturally the more interesting, it is easy to list a whole series of reasons why the intuitive approach to diagnosis is unsatisfactory. One reason consists in the large number of data required for diagnosis. When diagnosis is based on a questionnaire allowing only yes/no answers, experiments have shown that practically useful questionnaires must contain some hundreds of questions. Not all the questions, however, can be regarded as of equal importance. Add to this the fact that we usually have no *a priori* knowledge of the relationships between the symptoms and the underlying categories, nor about the importance of each individual symptom, and the complexity of this kind of problem becomes clear. In practical diagnosis one seems to single out the most important symptoms and to assess these, consciously or unconsciously, on a basis of previous experience, summarising the findings in order to make appropriate decisions. It is unlikely that a man could perform these operations with sufficient accuracy when large numbers of symptoms are in question. It is not surprising then that, in cases beyond the trivial, diagnosis resembles an art rather than a science.

3. Formalising the diagnostic problem

If formal procedures are to be used for diagnosis, we must first formulate the problem itself in the appropriate mathematical terms and work out procedures enabling us, on the basis of appropriate data, to classify a given situation into one or another category. For this purpose we shall use a geometrical framework for depicting the possible situations and categories. The symptoms will be depicted in n-dimensional Euclidean space with an axis for each symptom. If the system is such that each symptom can only have two values, e.g. present or absent, then every coordinate can only have two values which may as well be taken to be 1 (symptom present) and 0 (symptom absent). Such a geometrical model of the initial data is usually quite practical. All possible situations then appear as vertices of an n-dimensional cube.

As a simple illustration let us assume that the classification problem consists in classifying many situations into one of only two categories, A and B. The generality of our discussion is not really altered by this assumption since each of the two categories can subsequently be further subdivided. We assume, therefore, that some of the vertices of the n-dimensional cube belong to situations in category A and others to situations in category B. When this is so the problem of classification consists in finding a boundary (assuming that one exists) in the sympton space which separates the vertices of the cube belonging to category A from all other vertices. When this boundary has been located, then the category to which a particular situation belongs can be determined by establishing on which side of the boundary the corresponding point lies.

To make this formal procedure possible, it is important that the set of available symptoms should be adequate for classification. Without this, we may get vertices of the cube belonging simultaneously to both categories. Clearly in this case it would not be possible to define a suitable boundary. To avoid this it would be necessary to extend the symptom space by adding new symptoms to those already in use until classification becomes possible. The opposite situation may also occur, in which some of the symptoms used for classification are redundant in the sense that situations could be allotted to categories without them. In this case it is perfectly possible to establish the required boundary, but the classification procedure is unnecessarily complicated by superfluous features of the symptom space.

The naive solution to the problem would consist in listing all possible symptom complexes and noting to which category each of them belongs. However, when we remember that in cases of any complexity the number of symptoms is quite large, perhaps in the hundreds, it must be recognised that this kind of procedure, though possible in principle, is practically of no use since the number of possible symptom combinations, and consequently the number of vertices of the n-dimensional cube, is as high as 2^n. With $n = 100$, for example, which is quite common in practical investigations, the total number of possibilities is so high that even the fastest modern machine would take a million years to perform a classification. It follows that effective methods of classification cannot be based on this approach. Fortunately it has been found that establishing a boundary which is reliable in practice does not necessitate the study of all possible symptom complexes. Theoretical studies confirmed by practical investigation have shown that it is adequate to cover only a small fraction of all possible situations provided this fraction is sufficiently representative of the whole. For the solution of the problem it is sufficient to use a number of situations of the order of the logarithm of the total possible number. Thus in practical cases the number of situations needed to establish the boundary is not too high for the method to be practical on modern computers.

4. Computer learning as a classification procedure

There now exist formal procedures for 'training' a computer to undertake the task of classifying situations to different categories. Two possible approaches may be adopted and these differ notably from each other. In the first approach the trainee is provided with a procedure or algorithm by which the required task may be performed, whereas with the second approach training is done by example. Thus both human beings and machines are 'trained' to carry out arithmetic or logical operations by being provided with suitable algorithms. On the other hand procedures for tasks such as recognising letters or numerals in various type faces or written by hand are not imparted by explaining the structure of the symbols or by going into details of the recognition mechanism, with which the trainer is not himself familiar; these procedures are taught by example. When this approach is adopted it is important that the learner at the conclusion of training should be able to identify correctly new situations which have not been introduced during the training process itself.

Training by example is of great importance for the existence of many types of living organism. The method is used by animals teaching their young to find food and to avoid dangerous situations. Many habits and techniques which enable human beings to carry out a variety of tasks necessary for their continued existence are also acquired by observation or by analogy, without any explicitly defined algorithm.

Until recently automatic equipment, even when as advanced as digital computers, solved problems by means of clearly defined algorithms. We may say that in designing the structure of automata or by writing programs for computers, human beings teach the equipment to solve its particular tasks by way of the first approach described above. This approach, however, proves to be useless for tasks whose algorithmic solution is unknown to us even though they can be resolved intuitively. For example, a man can easily be taught to distinguish between a cat and a dog, to recognise his friends, or catch a ball in mid-air; but we do not know how to design a computer program to perform any of these actions.

A desire to expand the range of computer abilities has forced scientists and engineers to attempt to simulate the process of learning by example in computers. These attempts have already yielded positive results. The first successful effort was the design of automata capable of identifying visual patterns such as geometrical figures, letters, numbers and other symbols.

The task of pattern recognition may be described as follows: We have a cluster X containing a large number of objects $x_1, x_2, \ldots x_N$ and these objects belong to a relatively small number of known categories $X_1, X_2, \ldots X_m$. The recognition task is solved by an automaton if it can always (or at least with a known degree of reliability) classify each object x into a particular pattern category X_i. Thus suppose the cluster X represents all possible graphic symbols for numerals, then the recognition task consists in classifying each symbol into one of the ten categories 0, 1, 2, . . . 9. Alternatively the cluster X might contain a variety of geometric figures which have to be sorted into the categories of triangles, squares, circles, etc.

An automaton for pattern recognition must obviously contain an input device which can receive information on the pattern which is to be recognised. We call this the *receptor-field*. It may, for example, consist of a mosaic of photocells onto which the pattern to be recognised is projected. If the field consists of r cells, each capable of being in one of two states, then the number of possible configurations at the input is $N = 2^r$. Even with r quite small, N is so large that it is impossible in practice to store the data on every possible configuration in a computer memory.

The output device of the automaton must have m outputs. It is then possible to establish which category the pattern has been allotted to by determining which of the outputs has been activated.

So that the automaton may be trained it must have an adequate number of internal states $z = z_1, z_2, \ldots, z_m$, among which are states in which the automaton classifies the patterns in the required manner.

An automaton of this kind can be trained by providing a relatively limited number l of patterns whose categories are known, and arranging that the corresponding internal states z are made to provide the required output. It is found that by the time the automaton classifies correctly an adequate number l of patterns from the cluster X it is able to classify fairly satisfactorily all other patterns from the cluster as well. In more detail, the automaton is confronted with l objects selected arbitrarily from the cluster X. The trainer (which may be man or another automaton) indicates the correct category for each of the l objects. The trainer output is compared with the output of the automaton A and the internal states of A are modified to make its output correspond as closely as possible with that of the trainer. A theoretical limit to the number of matches needed to guarantee a given reliability for the trained automaton can be established but in practice a much smaller number is employed. At any stage the automaton may be confronted with a further arbitrary selection of patterns, and if the error rate is acceptable then training is considered complete. If not, a further stage of training can be introduced. This procedure is continued until the necessary degree of reliability is attained or until it is established that the particular automaton is incapable of being trained to recognise patterns on the basis of the characteristics presented to it.

The efficiency of the training procedure depends for the most part on three factors:

(1) the characteristics of the pattern which are presented to the receptor-field of the automaton;
(2) the number of possible transformations between the input configuration of the automaton and its output, that is to say the number of internal states;
(3) the detailed procedure of the training device.

The information about the pattern presented to the receptor-field must contain sufficient information for classification. The variety of internal states must be large enough to make it possible to train the automaton to tackle a sufficiently wide range of tasks, but not so large as to necessitate an excessive number of matches during the training phase. The procedure used by the training device must provide the highest possible recognition reliability foı a given number of training matches.

A very efficacious training procedure has been developed in the Instituute for Automation and Telemechanics at the Academy of Sciences, USSR, under my management (Vapnik, Lerner, and Chervonenkis, 1966). This enables a digital computer to develop, on the basis of a number of patterns belonging to various categories a simple formula representing a generalised characteristic for each pattern. When this formula has been obtained the classification task is solved quite simply by inserting into the formula the values of the characteristics typifying a particular pattern. It has proved possible to train computers for classification tasks to an extent which is sufficiently reliable for practical purposes.

5. Examples of the application of formal diagnostic techniques

Diagnostic methods based on the training of computers to classify complex situations are today being applied to the solution of many practical problems. I give here three problems whose successful solution has been made possible only by way of formal diagnostic methods. These problems have been tackled in my laboratory at various times. Our automated technique for diagnosis is now in widespread use and the problems I discuss have been selected to illustrate its extensive possibilities.

The first problem is taken from the area of petroleum geology. When drilling for oil it is necessary to distinguish between water-bearing and oil-bearing strata at the drilling site. Oil-bearing strata must be tapped to draw oil from them into storage reservoirs and it is essential to avoid tapping water-bearing strata by mistake to avoid the reservoirs being contaminated. Should this occur the drill hole has to be closed and drilling resumed elsewhere. This is an expensive matter, as is the second kind of error when an oil-bearing stratum is not tapped because of doubts that it might be water-bearing. Reliable diagnosis as between the two types of strata is thus of considerable importance. Geophysicists and petroleum geologists have therefore sought procedures leading to the reliable diagnosis of strata and have devised a standard complex of investigations at the drilling site. This complex includes electrical measurements, other measurements of artificial and natural radio-activity, etc. None of the individual measurements is capable by itself of defining the characteristics of the stratum, and diagnosis is in practice done intuitively on the basis of the totality of measurements in the complex. This is not at all a simple task and even the most experienced experts are mistaken in between 10 and 15% of all cases.

In 1963 we devised a formal technique for diagnosing oil-bearing strata. After training by our generalised method the computer produced a formula including more than ten parameters by which diagnosis can be performed. The diagnostic procedure can be carried through using this formula on the site without the need for elaborate calculating equipment. The error rate is one in 200 cases.

The second problem is concerned with the application of formal diagnosis to crime investigation. We undertook this task in cooperation with the Institute for Juridical Expertise. The problem involved devising a method for identification of handwriting. Criminological experts chose two people with similar handwritings and determined that the most easily confused character was 'b'. We therefore reproduced photographically a number of examples of this letter from texts written by the two people and trained the computer by our formal procedure to distinguish between the two handwritings. Thirty letters were used for training and the total study involved 160 letters from each person. The computer obtained correct diagnoses in over 90% of all cases, thus achieving a better reliability than the handwriting experts—in fact the computer was able to distinguish between the handwritings even better than the writers themselves.

My third example, and in my opinion the most interesting, involves medical diagnosis. We have studied a number of problems in the area of differential diagnosis and our incomplete results to date show that formal methods may be able to give considerably better results than human diagnosticians. I have no doubt that in the immediate future formal methods will be of immense help to physicians and will greatly increase the reliability of medical diagnosis. Moreover they will make it possible to achieve successful diagnosis at a much earlier stage than is possible with currently available techniques.

References

VAPNIK, V. N., LERNER, A. YA., and CHERVONENKIS, A. YA. (1966). Learning of a Machine to Pattern Recognition by the Method of Generalised Portrait, *Transactions of the third Congress of IFAC.*

Вапник, В. Н., Лернер, А. Я. (1963). Узнавание образов при помощи обобщенных портретов, Автоматика и телемеханика, т. XXIУ, No. 6.

Вапник В. Н., Лернер А. Я., Червоненкис А. Я. (1965). Системы обучения распознаванию образов при помощи обобщенных портретов, Иэвестия АН СССР, Техническая кибернетика, No. 1.

Book Review

Mathematical Linguistics in Eastern Europe, by Ferenc Kiefer, 1968; 180 pages. (Elsevier Pub. Co. Ltd., Barking, £6 0s. 0d.)

The most important result for linguistic theory of the first French-Russian machine translation algorithm was Kulagina's set-theoretic model of natural language. Kulagina's and other mathematical models of language are presented in varying detail by Dr. Kiefer (head of the machine translation group in the Computing Centre, Hungarian Academy of Sciences), examined critically and tested against five sensible criteria (nontrivial? formalised enough? linguistically relevant?, etc.). He restricts himself to theoretically important formal models but allows that the term 'mathematical linguistics' has wider application; in Eastern Europe, for historical reasons, it is often used simply as a synonym for 'modern linguistics'. (An illuminating definition, in the preface to 'Prague Studies in Mathematical Linguistics' Vol. 1, 1966, is: 'the quantitative analysis of language phenomena, the algebraic description of language systems, the theory of machine translation'.)

Šaumjan's 'applicational generative' model, as important a theoretical construct as Chomsky's or Lamb's, is well presented and examined. Though not trivial formally, Kulagina is found trivial when confronted with linguistic data. Other models presented are: Bierwisch on semantics and intonation; the author on semantics and the interrelation between emphasis and syntactic rules; Marcus on a formal model of the phoneme; contributions to the formal theory of transformations by the Czech, Čulik; Sgall's eclectic model (Prague School, Chomsky, Lamb) of a multilevel generative description of language; a dependency model of grammar by Filiatov (Leningrad).

I appreciated the presentation of the work of Šaumjan and Bierwisch and word of forthcoming contributions by the author on the order of elements in generative models (set-systems *v.* concatenation systems). I regretted lax proof-reading and an inexplicably high price. Some linguists may find the mathematical apparatus a bit fearsome; the documentalists and computer programmers mentioned on the jacket will have to be well up in current linguistic theory for parts of the book.

T. HAGAN (Ulster)