

Simple knowledge structures for perceptual classification expert systems

RICHARD BLOCH and SOREL BOSAN
College of William and Mary and Eastern State Hospital
Williamsburg, Virginia

A new approach is described for the development and structure of perceptual knowledge-based systems. This systematic method for acquiring perceptual knowledge for use in knowledge-based systems and for representing perceptual knowledge within the systems has revealed that two broad classes of perceptual activity can each be characterized by a single logical operator. The number of rules necessary to accomplish particular perceptual tasks can also be estimated.

An effort to develop techniques for automating the differential diagnosis of movement disorders led to an examination of methods for acquiring perceptual knowledge in a form that could be incorporated readily into a knowledge-based system. The potential for automated differential diagnosis of involuntary movement disorders such as Parkinson's disease, Huntington's chorea, tardive dyskinesia, or cerebral palsy revolved around the ability of a computer to acquire information about the parameters of movement and to be able to distinguish between normal and various forms of abnormal movement.

To acquire data about parameters of movement such as velocity, direction, and extent requires the ability to "see" the movement either through computer vision or through analog-to-digital conversion of voltage changes proportional to the movement. Once movement data is gathered, by an image processor, through psychophysiological measures such as EMG, or from accelerometers, the data must be reduced and interpreted if the movements are to be classified as normal or abnormal, or as one particular disordered movement type. The same data acquisition and interpretation steps are necessary for any perceptual classification problem, whether it involves sexing baby chickens, industrial quality-control monitoring for products such as ice cream or perfume, or voice recognition.

There are at least two approaches to automated classification: neural networks and expert systems. Neural networks model human learning at the neuronal level, making and strengthening connections as a function of the association of input characteristics with response outcomes. Successful implementations generally require many repeated exposures to the stimulus situations that must be discriminated. But with stimulus types that are relatively rare, this can be problematic.

Expert systems, on the other hand, incorporate the knowledge of human experts into knowledge structures that enable the computer to utilize the information to make

judgments. Knowledge acquisition for expert-system development has traditionally relied on verbal interviews, protocol analysis, and observation of selected experts. A knowledge engineer translates the experts' verbal descriptions of, and behaviors toward, decision-making variables into a knowledge-based structure, such as rules, and combines the rules in decision-making structures to emulate the expert's decisions. In knowledge domains where this approach can be used, the knowledge acquisition has proven to be a bottleneck for the development of expert systems (Hayes-Roth, Waterman, & Lenat, 1983). Current research to reduce the bottleneck is focused on the use of automated interviews (Gruber & Cohen, 1987) and on interfaces that gather knowledge directly from the expert, bypassing the translation of the expert's knowledge by the knowledge engineer (Kitto & Boose, 1987).

Despite current efforts to improve knowledge-acquisition methods, these methods are quite difficult to apply to perceptual domains. An alternative approach has been described and tested (Bosan, 1989), which utilizes interactive simulations to acquire knowledge. In this technique, experts are given a computer-controlled simulation of the stimulus situation of interest, with values of key features of the stimulus determined by the expert's use of interactive controls. Data obtained from an expert on each dimension represents that expert's judgment that, after appropriate adjustment to match a selected stimulus condition, the final adjusted simulation value represents the expert's estimation of that dimension for that stimulus condition. With repeated trials, the extent of adjustment variability can be indexed. Thus, for example, in adjusting the tempo of movement of a computer-controlled graphical display of a particular body part to match the stimulus condition Parkinsonian tremor, an expert neurologist finishes changing the tempo adjustment when it matches his or her knowledge of Parkinsonian tremor tempo. Repeated adjustments for tremor tempo would provide a basis for establishing a distribution of scores that could be standardized to provide their associated probabilities.

Rules are developed as a function of the number of dimensions and the number of conditions to be discrimi-

Sorel Bosan is now at Vanderbilt University, Nashville, TN. Correspondence may be addressed to Richard Bloch, Eastern State Hospital, Drawer A, Williamsburg, VA 23187.

nated. There are two ways to generate rules. One way is to create a single rule for each stimulus type. Thus, in the example above, there would be a single rule for Parkinsonian tremor, with clauses for every dimension, such as tremor tempo. Because each simulated dimension would be adjusted several times, the rule could have a range of true values that would be as inclusive or exclusive as desired. Each discriminated condition would have the same number of clauses in its single rule, and each clause would have the same structure. To extend the example begun above, in discriminating between Parkinsonian tremor (PT) and other movement conditions, each condition would have a single rule that had the following form:

```

IF: [value of dimension(tempo) in test stimulus] ≥
    [lower boundary of simulated condition(PT)
    dimension(tempo)] and ≤ [upper boundary of
    simulated condition(PT) dimension(tempo)]
AND [value of dimension(2) in test stimulus] ≥ [lower
    boundary of simulated condition(PT) dimen-
    sion(2)] and ≤ [upper boundary of simulated
    condition(PT) dimension(2)]
.
.
AND [value of dimension(N) in test stimulus] ≥ [lower
    boundary of simulated condition(PT) dimen-
    sion(N)] and ≤ [upper boundary of simulated
    condition(PT) dimension(N)]
THEN: [test stimulus] = [condition(PT)]
.
.
IF: [value of dimension(tempo) in test stimulus] ≥
    [lower boundary of simulated condition(2) dimen-
    sion(tempo)] and ≤ [upper boundary of simu-
    lated condition(2) dimension(tempo)]
.
.
AND [value of dimension(N) in test stimulus] ≥ [lower
    boundary of simulated condition(2)] and ≤ [upper
    boundary of simulated condition(2) dimension(N)]
AND ...
.
.
THEN: [test stimulus] = [condition(N)]

```

Although each rule has the same number and form of clauses, the rule structure has no particular relationship to the conditions being discriminated. The rule structure represents a general form for containing perceptual judgment ranges set by human experts. Also note that each clause is ANDed within a stimulus condition. For a rule to be true, each of the dimension clauses within the rule for the stimulus condition must be true. This is also a constant that applies across stimulus conditions and represents all classification/diagnostic problems.

The use of a range could be based on the statistical distribution of the experts' judgments. In the initial test of this method, the range was arbitrarily determined by the mean plus and minus two standard deviations. The range

could, however, be based on a variety of other statistical functions.

With this single rule-per-stimulus-condition format, a conflict resolution strategy may be required to make a final decision. If, in the example above, more than one rule were true, this strategy would be used to determine the output. Furthermore, this strategy could be represented as an additional rule. This rule could decide that if a test stimulus yields more than one stimulus condition to be true, or if no stimulus condition rules are found to be true, then the test stimulus condition is unknown or is a stimulus condition other than those simulated.

With a single rule for each stimulus condition, the number of rules equals the number of discriminated conditions plus the final conflict resolution rule. Given the capabilities of today's rulebase shells, the limiting factor in developing rulebases that classify dozens, or even hundreds, of perceptual stimuli becomes the time required for expert input. If, as with neural networks, this method of acquiring and organizing knowledge is scalable, the multiple smaller perceptual rulebases might be able to be combined without requiring expert reexamination of perceptual simulations.

An alternative to single rules per stimulus condition is to break each clause in the single rule example into its own rule. This would mean that for every stimulus condition being classified or discriminated, there would be the same number of rules as there are dimensions being adjusted. Using the same example of Parkinsonian tremor, this approach would produce a rule base in the form of:

```

IF: [value of dimension(tempo) in test stimulus] ≥
    [lower boundary of simulated condition(PT)
    dimension(tempo)] and ≤ [upper boundary of
    simulated condition(PT) dimension(tempo)]
THEN: [test stimulus dimension(tempo)] = [PT dimen-
    sion (tempo)]
IF: [value of dimension(2) in test stimulus] ≥ [lower
    boundary of simulated condition(PT) dimen-
    sion(2)] and ≤ [upper boundary of simulated
    condition(PT) dimension(2)]
THEN: [test stimulus dimension(2)] = [PT dimension(2)]
.
.
IF: [value of dimension(N) in test stimulus] ≥ [lower
    boundary of simulated condition(condN) dimen-
    sion(N)] and ≤ [upper boundary of simulated
    condition(condN) dimension(N)]
THEN: [test stimulus dimension(N)] = [condition(condN)
    dimension(N)]

```

Again, all the rules within the rulebase would have the same basic form, even though they were relevant to different dimensions of different stimulus conditions. Additional rules would be necessary to make the expert system's final judgment. For each stimulus condition, there would be one rule that combined all the dimensions with a logical AND such as:

IF: [test stimulus dimension(tempo)] = [PT dimension(tempo)]
 AND [test stimulus dimension(2)] = [PT dimension(2)]
 AND ...
 .
 .
 AND [test stimulus dimension(N)] = [PT dimension(N)]
 THEN: [test stimulus] = [PT]

These combination rules would all use the logical AND for each dimension, independent of what was being classified or diagnosed. In short, the logical AND, whether within a single rule per condition, or between single rules per condition \times dimension combination, characterizes all perceptual classification and diagnostic problems.

Again there is one additional conflict resolution strategy rule that is needed to handle the situation in which none of the rules combining dimensions for a stimulus condition are found true or more than one is found true. Thus, when utilizing single rules per dimension, the total number of rules required to classify N stimulus conditions becomes a function of the formula: no. of rules = (no. of stimulus conditions) \times (no. of dimensions) + 1.

Just as the logical AND applies to classification and diagnosis problems in using the interactive simulation method to acquire knowledge for perceptual expert systems, it is proposed that the logical OR applies to perceptual detection problems. In this context, detection is composed of a search strategy for scanning in space or time for a stimulus event. It is the search rather than the event recognition that is central to the detection process. For example, in scanning a surface for an imperfection, gaze is moved from one location to the next, whereas in listening for a sound, both spatial orientation to potential sources and repeated samplings over time occur. The scanning process essentially ends with the detection of a stimulus event, such as would be signaled by an orienting response.

The use of interactive simulations to acquire knowledge for a detection expert system is not as well understood as their use to acquire knowledge for classification problems. It can be proposed that the dimensions adjusted by the expert involve parameters or dimensions of the scanning process such as the order in which the segments of the whole scan area are scanned; the size of the individual scan segments; the rate of scanning each segment; and the degree of change from background required to stop scanning and cause an assessment as to whether a stimulus is the sought-after condition. In sum, these dimensions amount to a detection strategy.

The rules generated in detection problems are also suggested to be simple and systematically structured. Just as with classification problems, each complete search activity within a sensory modality can have a single rule, such as:

IF: search segment(1) = stimulus change
 OR search segment(2) = stimulus change

OR ...
 .
 .
 OR search segment(N) = stimulus change
 THEN: stimulus change detected

Another example would be:

IF: search segment size(1) = stimulus change
 OR search segment size(2) = stimulus change

THEN: stimulus change detected

Again, it should be noted that in the class of detection activities, logical ORs rather than ANDs link the rule clauses. The number of clauses within a rule is a function of the number of search segments. If the assessment of the stimulus change is not considered part of the detection process, each sensory modality may require a single rule. Thus, the number of rules may be a function of the number of sensory modalities scanned. If more than a single modality is scanned, the rule structure would become:

IF: search segment(1) modality(1) = stimulus change
 OR search segment(2) modality(1) = stimulus change

OR search segment(N) modality(1) = stimulus change
 THEN: stimulus change modality(1) detected

Each additional modality would have a similar rule. The combination rule to decide if an event was detected could again use the logical OR to produce a decision that an event was either detected or not detected; for example:

IF: stimulus change modality(1) detected
 OR stimulus change modality(2) detected
 OR ...
 THEN: stimulus change detected.

Returning to the original goal of developing a knowledge-based system capable of discriminating between different movement disorders, it becomes clear that most perceptual tasks begin with an embedded detection problem. In order to classify movement disorders, movement must first be detected. Once it is detected, the classification process begins. The design of an expert system capable of this sequence of activities would, thus, require a scanning strategy based on one or more experts' scanning strategies, and classification rules that are tested once the scanning process produces a stimulus detection. The classification process would determine if the detected stimulus was, for example, a movement or not, and then, if it was judged a movement, would classify the movement type.

In this way, detection and classification rules are the basic building blocks of perceptual behavior, and they enable perceptual behavior to be modeled. Complex perceptual expert systems become conceivable as a combination of relatively simple detection and classification problems with consistent rule structures and various se-

quences of OR and AND decision structures. Through the building of sequences of detection and classification tasks, the potential is created to develop knowledge-based systems that could detect and classify various objects; for example:

1. Barriers in a robotics environment could be detected and classified in terms of what types of barrier or what types of avoidance behavior should follow.

2. Movement in a security monitoring environment could be detected and classified as resulting from human or other causes.

3. Manufactured objects on an assembly line could be examined for defects; once detected, they would be classified as to whether they were correctable or noncorrectable.

In summary, the first systematic method for acquiring perceptual knowledge for use in knowledge-based systems, and for representing perceptual knowledge within knowledge-based systems, has pointed toward surprisingly simple and consistent rule structures. The process has revealed that two broad classes of perceptual activity can each be characterized by a single logical operator. In the case of classification and diagnostic activities, the logical operator is an AND, whereas in detection activities,

the operator is an OR. It has been proposed that complex perceptual activities can be broken down into a sequence of steps, each step utilizing one of these activity classes, and one of the two operators. The number of rules necessary to accomplish particular perceptual tasks can also be estimated, with classification and diagnostic activities being quite efficient in the number of rules necessary for rudimentary discriminations. This represents a new approach to the development and structure of perceptual knowledge-based systems and holds promise for a wide variety of industrial and medical problems.

REFERENCES

- BOSAN, S. (1989). *Interactive simulation for knowledge acquisition*. Unpublished MA thesis, College of William & Mary, Williamsburg, VA.
- GRUBER, T. R., & COHEN, P. R. (1987). Design for acquisition: Principles of knowledge-system design to facilitate knowledge acquisition. *International Journal of Man-Machine Studies*, **26**, 143-159.
- HAYES-ROTH, F., WATERMAN, A. D., & LENAT, D. B. (1983). *Building expert systems*. Reading, MA: Addison-Wesley.
- KITTO, C. M., & BOOSE, J. H. (1987). Choosing knowledge acquisition strategies for application tasks. *IEEE Proceedings*, **8**, 96-103.