

ANALYSIS OF AN EXTENDED CONCEPT-LEARNING TASK

Richard M. Young, Gordon D. Plotkin
(University of Edinburgh)

and

Reinhard F. Linz
(University of Bonn)

Bruner, Goodnow and Austin (1956) describe a "focussing strategy" for learning simple conjunctive concepts in a situation where the learner is shown a sequence of instances, one at a time, each being an example or non-example of the concept to be learned. We consider the extension of this strategy to concepts in which the individual features are hierarchically structured, as is the case, for example, for the materials used by Winston (1975). The essence of the extended strategy is to represent an hypothesis by a *pair* of nodes in the hierarchy, instead of just a single node. The "lower", more specialised node indicates the set of instances that have already been inferred to be examples of the concept. The "upper", more abstract node indicates an inferred limit on how general the concept can be. Thus the nodes act as bounds between which the concept must lie, and represent respectively sufficient and necessary conditions for an instance to be an example of the concept.

We generalise this notion of an hierarchically-structured feature by defining a *description space* (d-space) as an upper semilattice $(D, \leq, \vee)$, with a top element, T , satisfying a certain finiteness condition. We show that the Cartesian product of a number of these d-spaces is itself a d-space. This provides us with a simple, uniform notation for referring to concepts, instances and hypotheses.

We describe a simple nondeterministic algorithm which implements the "focussing strategy" for such a conjunctive description space, and discuss aspects of its behaviour on the concept learning task. For example:

1. Due to the generality of the d-space idea, the algorithm deals uniformly with features which in Winston's approach have to be treated as special cases. Predicates, relations ($A \text{ SUPPORTS } B$), multi-valued dimensions (ORIENTATION) and tree-structured dimensions (SHAPE) all give rise to examples of a d-space.
2. Given appropriate assumptions, it can be proved (a) that the hypothesis generated by the algorithm is always consistent with the data, *i.e.* it includes all the examples and excludes all the non-examples it has so far been shown, and (b) that given sufficient data, the algorithm correctly learns the concept.
3. According to the algorithm, back-up occurs when the relation between the lower and upper bounds is violated. The lower node of a hypothesis is affected only by positive examples, and the upper node only by non-examples. This clarifies the role of negative information and "near misses" in the concept learning. In particular, it can be shown that . . .
4. . . . Given a corpus of instances, the outcome of the learning is independent of the order in which they are presented. Moreover . . .
5. . . . If all the examples are presented before the non-examples, then the learning is free of errors, *i.e.* the algorithm never needs to back up.

In the light of this analysis, we argue that Winston's program is itself an attempt to use a focussing strategy to learn conjunctive concepts with hierarchically-structured features, and we discuss various shortcomings and issues not dealt with in Winston's treatment.

REFERENCES

- Bruner, J. S., Goodnow, J. J. & Austin, G. A. *A Study of Thinking*. Wiley, 1956.
- Winston, P. H. Learning structural descriptions from examples. In P. H. Winston (Ed.), *The Psychology of Computer Vision*. McGraw-Hill, 1975.