

Multiple Convergence: An Approach to Disjunctive Concept Acquisition *

Kenneth S. Murray

Department of Computer Sciences
The University of Texas at Austin
Austin, Texas 78712

ABSTRACT

Multiple convergence is proposed as a method for acquiring disjunctive concept descriptions. Disjunctive descriptions are necessary when the concept representation language is insufficiently expressive to satisfy the completeness and consistency requirements of inductive learning with a single conjunction of generalized features. Multiple convergence overcomes this insufficiency by allowing the disjuncts of a complex concept to be acquired independently. By summarizing correlations among features in the training data, disjunctive concepts can provide rich extensions to the representation language which may enhance subsequent learning. This paper presents the benefits of disjunctive concept descriptions and advocates multiple convergence as an approach to their acquisition. Multiple convergence has been implemented in the learning system HYDRA, and a detailed example of its execution is presented.

I Introduction

The learning problem addressed is concept acquisition from examples, as formulated in [MITC82]. This entails developing a concept description to summarize training objects, each classified as a positive or negative instance of the target concept. The summary description is constrained to admit every positive instance and to reject all negative training. These are called the *completeness* and *consistency* requirements [MICH83]. The concern of most machine learning research in inductive concept acquisition has been the determination of *characteristic descriptions*, which represent concepts by summarizing the properties that hold true for all instances of the concept [DIET83]. Characteristic descriptions are typically encoded as a single conjunction of maximally specific features.

A. Classical Concepts

Classical concept descriptions are composed of features that are individually necessary, and jointly sufficient, for classification of objects as concept instances [SMIT81]. Because of their single-conjunction form, the characteristic descriptions produced by most existing concept acquisition systems resemble classical concept descriptions.

Unfortunately, a single conjunction of necessary and sufficient features is not suitable to represent many concepts in natural domains. Figure 1 illustrates this with a simple example of learning about *trees*. After two positive training instances, a classical concept description is adequate. However, when a negative instance is encountered, the existing classical description is found overly-general and must be refined into a disjunction. No single conjunction can satisfy both the completeness and consistency requirements for the given training and language.


Learning systems that produce only classical concept descriptions reflect the assumption that concepts are *independently separable*. A concept is independently separable when the independent determination of admissible values for each feature is sufficient to distinguish all positive instances of the concept from all negative instances. Concept acquisition under this assumption can be characterized by the following two steps:

- 1) for each attribute, find the set of values valid for the target concept
- 2) take the cartesian product of all sets defined by step 1

This approach fails when the attributes of the target concept interact, and the values allowed for an attribute in one context may not be valid in another. Unless all necessary featural correlations are present in the representation language, a system limited to classical concept descriptions cannot learn these concepts.

CONCEPT TO BE LEARNED: *tree*

REPRESENTATION LANGUAGE:



TRAINING:

- instance 1: *positive*; [leaf, apple]
characteristic summary: [leaf, apple]
instance 2: *positive*; [needle, cone]
characteristic summary: [foliage, seed]
instance 3: *negative*; [needle, pear]
characteristic summary: [leaf, apple] OR [needle, cone]

Figure 1: Learning About Trees

* Support for this research was provided in part by the Army Research Office under grant number ARO DAAG29-84-K-0060.

B. Disjunctive Concepts

Disjunctive concept descriptions are necessary when the representation language is insufficiently expressive to produce a complete and consistent classical description. This situation arises when either of two anomalies occurs:

- 1) *incompleteness condition*: a positive instance is encountered, but the concept description cannot be generalized to admit it without also admitting negative training
- 2) *inconsistency condition*: a negative instance is encountered, but the concept description cannot be specialized to reject it without also rejecting positive training

A disjunctive concept description resolves these dilemmas by relaxing the completeness requirement for the individual disjuncts. Each disjunct summarizes some subset of positive training while remaining consistent with all negative training. Collectively, the disjuncts satisfy both completeness and consistency. Identification using disjunctive concept descriptions requires classifying an object as an example of the concept when any disjunct admits it.

II Multiple Convergence

Multiple convergence enables acquisition of characteristic descriptions for disjunctive concepts. It differs from existing strategies in the way concepts are represented and how disjuncts are created and updated, allowing each disjunct to emerge as an independent concept representing a specialization of the target concept. Multiple convergence has been implemented in the learning system HYDRA.

A. Representing Disjunctive Concepts

Although the objective is to produce a complete and consistent characteristic description of the target concept, multiple convergence maintains a *discriminant description* as well as a characteristic description of each concept it learns. The discriminant description summarizes negative training and defines criteria that all instances of the concept must *necessarily* satisfy. The characteristic description summarizes positive training and represents hypothesized *sufficient* criteria. As with Focussing [BUND85] and Candidate Elimination [MITC78], concept formation proceeds as a bi-directional convergence: the discriminant description progresses from general to specific with negative training while the characteristic description generalizes with positive training.

Disjuncts formed in the discriminant description are used to define the disjuncts of the characteristic description. Intuitively, the discriminant description identifies regions of the feature space that can accommodate a classical description. Each disjunct of the discriminant defines such a region. A characteristic disjunct (i.e. a single conjunction of generalized features) is developed to describe the positive instances within that region.

This process applied to the example from Figure 1 produces discriminant description [*foliage, seed*] for the first two training instances. After encountering the negative instance, the discriminant description is specialized into three disjuncts: [*foliage, apple*], [*foliage, cone*], and [*leaf, seed*]. Each of these defines a region of the feature space that can currently accommodate a classical description. For each region, a characteristic disjunct is developed to represent the positive training within that region, producing the characteristic disjuncts [*leaf, apple*] and [*needle, cone*], as shown in Figure 1.

Each associated pair of disjuncts from the discriminant and characteristic descriptions defines a *version space* [MITC78] for a hypothesized specialization of the target concept. As the version space of each disjunct converges, it defines a classical description of a *generalized exemplar*. The target concept is represented by the disjunction of the surviving set of generalized exemplars. In the example of Figure 1, two generalized exemplars representing deciduous and coniferous trees (e.g. [*leaf, fruit*] and [*needle, cone*]) will survive the convergence induced by exhaustive training.

B. Updating Disjunctive Concepts

Positive training instances are allocated to every disjunct that can be consistently generalized to admit them (as in [IBA79]). This is motivated from a desire to remain insensitive to training order. Rather than make arbitrary decisions that may require subsequent backtracking, all alternative hypotheses are maintained.

When the existing discriminant concept description is found to admit a negative training instance, it is replaced with a minimal-specialization. This involves specializing each disjunct only to the extent necessary to reject the training instance, and then retaining only those new disjuncts not more specific than some other consistent discriminant disjunct. Typically, this will introduce one or more new disjuncts into the discriminant description, and unlike other similar techniques [MITC78, IBA79, BUND85], disjuncts that fail to cover any existing positive instances are retained to avoid subsequent backtracking. Disjuncts of the characteristic description that admit the negative training instance must also be specialized. Furthermore, the characteristic description may need to be extended with new disjuncts to summarize the positive training admitted by new discriminant disjuncts.

Specializing and extending the characteristic description are tasks that have required most inductive learning systems to reprocess all prior positive training (e.g. STAR [MICH83], and extensions to Candidate Elimination [MITC78] and Focussing [BUND85]). Under multiple convergence, the reprocessing of prior training can be constrained to consider only a portion of the past positive training instances. This is achieved by indexing positive instances under the disjuncts that admit them [IBA79].

Each discriminant disjunct, d , found to admit a negative training instance is specialized to a set S of one or more new disjuncts. The reprocessing of past positive training to develop new characteristic summaries can be limited to those instances admitted by and indexed under d . This follows since each new disjunct s in S is a specialization of d , so the instances s admits must be a subset of the instances admitted by d . Therefore only the instances admitted by d need be considered when developing the new characteristic summary for s .

Reprocessing all the instances of an invalidated disjunct is avoided by retaining a trace of the prior characteristic summaries along with the positive training instances. A new positive training instance and the prior characteristic generalization are indexed under each new characteristic summary. As the disjunct "fills up" with positive training, an *exemplar generalization hierarchy* emerges. The exemplar generalization hierarchy is always rooted at the current characteristic summary of the disjunct, and its leaves are the positive instances admitted by the disjunct.

This data structure can be used to enhance reprocessing required by an inconsistency condition. For each s in S , a new characteristic summary is developed during an all-paths traversal of the exemplar generalization hierarchy associated with d . Each path can terminate when a node is either admitted by or disjoint from the new disjunct s . When a node of the exemplar generalization graph is admitted by s , then the node itself can be reprocessed as a positive instance, and all the actual instances indexed under it can be ignored. Similarly, if the node is completely disjoint from s , then all instances indexed under the generalization must be irrelevant.

The multiple convergence approach to learning disjunctive concepts has been implemented in the learning system HYDRA. An overview of the concept acquisition algorithm of HYDRA is presented in Figure 2, and the approach is illustrated by an example of learning a necessarily disjunctive concept in Figure 3.

```

REPEAT
  Await the next training instance,  $TI$ ;
  IF  $TI$  is positive
    THEN
      for each discriminant disjunct  $d$  that admits  $TI$  do
        minimally generalize the characteristic summary of  $d$ 
          to admit  $TI$ ;
  IF  $TI$  is negative
    THEN
      BEGIN
        minimally specialize the discriminant description
          to reject  $TI$ ;
        for each new discriminant disjunct  $d'$  do
          reprocess prior positive training admitted by  $d'$ 
            to produce a characteristic summary for  $d'$ 
      END;
  Display the discriminant and characteristic descriptions;
UNTIL the teacher is satisfied.
```

Figure 2: The Learning Algorithm of HYDRA

III Extending the Representation Language

Disjunctive concepts are inherently taxonomic and introduce new generalization hierarchies rooted at the target concepts. Each node of the new generalization hierarchy is a list of attribute-value pairs representing an instance or generalization in the feature space of the concept. Since every attribute is defined by a pre-existing generalization hierarchy, the new hierarchy of a disjunctive concept can be viewed as an *orthogonal generalization hierarchy*. The nodes of the new hierarchy represent useful correlations of features. The links define context-sensitive generalizations of the features. Upon the conclusion of a training session, dialogue with the teacher enables naming the useful nodes in the hierarchy. The named nodes are then elevated to the status of concepts. This process is illustrated by the example of Figure 3. The resulting orthogonal generalization hierarchy for this example is presented in Figure 4.

Shifting the generalization biases of representation languages used for inductive concept formation has been recognized as an important research topic, but previously only techniques to weaken the existing biases have been developed, [BUND85, UTGO86]. This involves focusing the *internal disjunction* provided by a generalization hierarchy by inserting a new node to segregate valid from invalid values of an attribute. The new generalization hierarchies introduced by disjunctive concepts represent an entirely new source of bias for generalization.

The new generalization hierarchies introduced by disjunctive concepts can enhance subsequent learning tasks. For example, consider learning about the concept *Italian-Exports* after having already defined *vehicle*. If provided with positive training examples [*pedals, handlebars*] and [*engine, handlebars*], the new generalization hierarchy enables immediate generalization to *vehicle* (i.e. Italy exports vehicles). Without prior knowledge of vehicles, this training would lead to the more conservative (and incorrect) generalization [*power, handlebars*].

IV Conclusion

Introducing disjunction into our concept descriptions broadens the class of concepts our systems can learn. Multiple convergence has the advantage of producing classical concepts when appropriate, but is also able to describe a concept as a disjunction of generalized exemplars. This becomes necessary when the representation language is insufficiently expressive to produce a single-conjunct summary that is both complete and consistent.

Many existing learning systems cannot learn disjunctive concepts because they assume the target concept is, independently separable. Multiple convergence puts only tentative faith in this assumption, allowing generalization operators to proceed as if it were so, but enabling relatively graceful recovery when the assumption fails.

Often the generalized exemplars of a disjunctive target concept themselves represent worthwhile concepts. For example, while learning about the disjunctive concept *vehicle*, multiple convergence also defines the concepts *bicycle*, *sailboat*, *motorized vehicle*, *car*, *motorcycle*, and *ship*. Therefore, multiple convergence emerges as an approach to learning multiple concepts simultaneously. By defining a generalization partial-ordering over the new concepts, disjunctive concepts enable rich extensions to the representation language.

ACKNOWLEDGEMENTS

I would like to thank Joe Ross, Ray Bareiss and especially Bruce Porter for their many insightful and encouraging comments.

REFERENCES

- [BUND85] Bundy, A., Silver, B., and Plummer, D. An Analytical Comparison of Some Rule-Learning Programs. *Artificial Intelligence*, Vol. 27(2) (1985).
- [DIET83] Dietterich, T. and Michalski, R. A Comparative Review of Selected Methods for Learning. In *Machine Learning*, Tioga Publishing, 1983.
- [IBA79] Iba, G. Learning Disjunctive Concepts From Examples. A.I. Memo 548, MIT AI Lab, 1979.
- [MICH83] Michalski, R. A Theory and Methodology of Inductive Learning. In *Machine Learning*, Michalski, R., Carbonell, J., and Mitchell, T. (eds.), Tioga Publishing, 1983.
- [MITC78] Mitchell, T.M. *Version Spaces: An Approach to Concept Learning*, PhD Dissertation, Computer Science Department, Stanford University. (TR STAN-CS-78-711), 1978.
- [MITC82] Mitchell, T.M. Generalization as Search. *Artificial Intelligence*, Volume 18 (1982), 203-226.
- [SMIT81] Smith, E. and Medin, D. *Categories and Concepts*, Cambridge: Harvard University Press, 1981.
- [UTGO86] Utgoff, P.E. Shift of Bias for Inductive Concept Learning. In *Machine Learning*, Michalski, R., Carbonell, J., and Mitchell, T. (eds.), Morgan Kaufmann, Vol. 2, 1986.

CONCEPT TO BE LEARNED: *vehicle*

REPRESENTATION LANGUAGE:



TRAINING:

instance 1: *positive*; [e, h]
 discriminant summary: {[po, st]}
 characteristic summary: {[e, h]}
 instance 2: *positive*; [p, h]
 discriminant summary: {[po, st]}
 characteristic summary: {[po, h]}
 instance 3: *negative*; [s, w]
 discriminant: {[e, st], [p, st], [po, h], [po, r]}
 characteristic: {[e, h], [p, h], [po, h]}
 instance 4: *positive*; [s, r]
 discriminant: {[e, st], [p, st], [po, h], [po, r]}
 characteristic: {[e, h], [p, h], [po, h], [s, r]}
 instance 5: *positive*; [e, r]
 discriminant: {[e, st], [p, st], [po, h], [po, r]}
 characteristic: {[e, st], [p, h], [po, h], [po, r]}
 instance 6: *negative*; [p, r]
 discriminant: {[e, st], [p, w], [po, h], [s, r]}
 characteristic summary: {[e, st], [po, h], [s, r]}
 instance 7: *negative*; [p, w]
 discriminant summary: {[e, st], [po, h], [s, r]}
 characteristic summary: {[e, st], [po, h], [s, r]}
 instance 8: *negative*; [s, h]
 discriminant summary: {[e, st], [p, h], [s, r]}
 characteristic summary: {[e, st], [p, h], [s, r]}
 instance 9: *done*
 Concept *vehicle* has 3 exemplars: {[p, h], [s, r], [e, st]}
 would you like to name [p, h]? (name, skip, quit) *bicycle*
 would you like to name [s, r]? (name, skip, quit) *sailboat*
 would you like to name [e, st]? (name, skip, quit) *motorized_vehicle*
 Concept *motorized_vehicle* has 3 exemplars: {[e, h], [e, r], [e, w]}
 would you like to name [e, h]? (name, skip, quit) *motorcycle*
 would you like to name [e, r]? (name, skip, quit) *ship*
 would you like to name [e, w]? (name, skip, quit) *car*
 < vehicle training session lasted 3933 ms >

NEW REPRESENTATION LANGUAGE: SEE FIGURE 4.

Figure 3: HYDRA Learning About Vehicles




Figure 4: Orthogonal Generalization Hierarchy