

KNOWLEDGE ACQUISITION FROM STRUCTURAL DESCRIPTIONS¹

Frederick Hayes-Roth
The RAND Corporation
Santa Monica, CA 90406

and

John McDermott
Carnegie-Mellon University
Pittsburgh, PA 15213

Abstract. The representation of concepts and antecedent-consequent productions is discussed and a method for inducing knowledge by abstracting such representations from a sequence of training examples is described. The proposed learning method, interference matching, induces abstractions by finding relational properties common to two or more exemplars.

1. INTRODUCTION

A number of distinct paradigms for studying learning machines have emerged during the last twenty years. Though each differs from the others in a variety of ways, the three differences which most clearly demark each paradigm are (1) the types of knowledge which can be acquired, (2) the way in which this knowledge is represented, and (3) the type of learning algorithm used. The learning machine which we will describe in this paper acquires concepts representable as conjunctive forms of the predicate calculus and behaviors representable as productions (antecedent-consequent pairs of such conjunctive forms); these concepts and behavior rules are inferred from sequentially presented pairs of examples by an algorithm that is provably effective for a wide variety of problems.

Learning is viewed here as a continual process of knowledge expansion, that is, as the acquisition, in adaption to training experiences, of higher-order, more complex, and more elaborate knowledge structures. One's knowledge at any point in time includes those concepts and productions innately provided or previously learned. The concepts are pattern templates; events which match a concept are recognized as belonging to the class delimited by that concept. The productions are pairs of concepts; one of the concepts functions as a recognizer, the other specifies the form of an associated action. A production is interpreted as a behavior generator in the sense that (in some computing environment with an appropriate control structure) the detection of a condition in the environment which matches the antecedent causes the consequent component to be instantiated and then evoked. Here both the antecedent and the consequent are templates; the antecedent determines whether the production is to be executed, and if so, what specific constants in the description of the event being attended to are to be bound to variables in the consequent.

Within this framework, the machine learning problem with which we are concerned can be stated in the following way: Given a collection of concepts and productions constituting what is known at some time and a way of describing events in terms of their structure,

Example 1



Example 2



Example 3




Figure 1

construct a machine which is able to induce additional concepts or productions from training data. To make our treatment of this problem more concrete, we will use the simplest of the concept formation tasks attempted by our machine as an example throughout the paper. The task is to find what the three exemplars in Figure 1 have in common. Our program induces the following abstraction:

There are three objects, including a small circle and a small square. The square is above the circle. The third object is large.

This paper is divided into four sections. In the next section we discuss in general a way of describing events which facilitates finding what two or more events have in common and a matching algorithm which can be used to find these abstractions. Then we locate SPROUTER, our concept and production inducing program, within the broader context of our work. The third section describes SPROUTER's interference matching (induction) algorithm in some detail; we indicate here more specifically how SPROUTER makes use of structural representations of events to acquire and store knowledge. In the final section we conclude with a brief consideration of the strengths and weaknesses of SPROUTER and directions for future research.

II. STRUCTURAL REPRESENTATIONS

The problem which we are addressing is simply described: Design a program which can infer concepts and productions from illustrative instances. The method we employ is correspondingly straightforward: Extract commonalities from the examples and attenuate their differences. Such an approach is like Galton's very primitive "composite photograph theory" of concept learning [8] and the "positive focusing strategy" for conjunctive concept learning first studied by Bruner, et al. [3]. While Galton's contribution was simply to propose that unknown patterns could be inferred by overlaying homologous memory representations of related examples (as if one were forming a composite of many photographs of the same subject), Bruner and his colleagues showed how such a process could in fact be realized. Each presented object (exemplar) is described as a conjunction of specific feature values. To find the template which is matched by all of the presented objects, a feature vector containing only those features common to all of the

¹ This research was performed while both authors were at Carnegie-Mellon University. It was supported in part by the Defense Advanced Research Projects Agency (F44620-73-C-0074) and monitored by the Air Force Office of Scientific Research. A more detailed version of this research report can be obtained from the Computer Science Department, Carnegie-Mellon University, Pittsburgh, PA 15213.

exemplars is generated. This feature vector is the concept. Since that seminal work, many computer scientists have produced increasingly practical and sophisticated feature-value concept learners based on related techniques [4, 10, 16, 20, 21, 29].

Extending such learning models so that they can induce general (relational) classification and behavior rules is the goal of our work. In focusing on methods for generating relational abstractions which make possible the recognition of complex events, we encounter three problems not encountered in previous work. First we must develop a formal scheme for describing complex events which facilitates the generation of abstractions. Second, given descriptions of two examples of the same concept or production, we must develop a method for comparing them so that their commonalities can be identified. Third, it is necessary to develop a way of storing the discovered abstractions to facilitate their subsequent use in either of two ways: they may be used as templates for classification and behavior generation, or they may be used as knowledge representations whose precision may later be improved by learning if new instances of the same concept or production are provided. These problems are referred to below as the description problem, the comparison problem, and the storage problem. Each is considered in more detail in the subsequent paragraphs.

The description problem entails providing a symbolic representation of each exemplar satisfying two demands. First, those attributes of the exemplar which are salient and potentially critical must be reflected in its description to insure that the classification rule induced will be sufficiently discriminating. Note that since an exemplar may be composed of many objects, the description must distinguish each object and indicate clearly how it relates to the others. Second, the descriptions should facilitate the identification of commonalities among the exemplars so that the abstraction being sought can be found quickly. Since each object may exhibit a variety of characteristics and participate in numerous relationships with other objects, finding commonalities between two or more examples will necessitate search. A representational scheme which helps direct this search is almost essential.

The method of description we employ is built on three central concepts, the property, the case frame, and the parameter. A property is a feature or characteristic of an object. For example, SQUARE and SMALL name two properties of small squares; the properties ABOVE and BELOW are used in our work to describe objects which are above or below others in pictorial displays. To define the relationship of one object being above another, a case frame of the sort {ABOVE, BELOW} is used. In general, case frames are sets of properties which are semantically related in some exogenously determined manner. To produce descriptions of objects, events, or behaviors, case frames are parameterized (instantiated); that is, a name is given to each object in the event being described and this name is associated with each property of the object. Parameterized case frames are called case relations. For example, if b is the name of a square above a circle named c, this might be described by the following set of case relations: {{SQUARE: b}, {CIRCLE: c}, {ABOVE: b, BELOW: c}}. Such a set of case relations interpreted as a conjunction of valid propositions is called a parameterized structural representation or PSR [9, 12, 13]. In this example, {b, c} is the parameter set of the PSR.

A structural description of the first two exemplars in the concept formation task discussed in the introduction is given below.

```
E1:
{ {TRIANGLE:a, SQUARE:b, CIRCLE:c},
  {LARGE:a, SMALL:b, SMALL:c},
  {INNER:b, OUTER:a},
  {ABOVE:a, ABOVE:b, BELOW:c},
  {SAME!SIZE:b, SAME!SIZE:c}}
```

```
E2:
{ {SQUARE:d, TRIANGLE:e, CIRCLE:f},
  {SMALL:d, LARGE:e, SMALL:f},
  {INNER:f, OUTER:e},
  {ABOVE:d, BELOW:e, BELOW:f},
  {SAME!SIZE:d, SAME!SIZE:f}}
```

The description of E1 asserts that there is an event composed of three objects, named a, b, and c; that the object labeled a has the properties of a triangle, of a large object, and of containing the object labeled b; and so on.

PSRs provide a solution to the storage problem as well as to the description problem; that is, they can be used in storing discovered abstractions. In the case of descriptions, parameter symbols are chosen to name each object so that if the same object is part of more than one case relation, it is referred to in a consistent way. If one alters the interpretation so that each distinct parameter is considered as an unbound variable, the PSR can be considered a template for concept identification. Such templates have been used by several researchers [1, 9, 12, 13, 15, 28] to specify what properties an object must have in order to satisfy membership in a pattern class. While the parameters in a description can be thought of as being existentially quantified, those in a PSR used as a template should be thought of as being universally quantified. When used as a template for pattern classification, the PSR is compared with an event (an existentially quantified PSR). If a mapping from the event to the template can be found which preserves the parameter bindings in the event description and which makes each case relation of the template true, the event is said to match the template.

In addition to their role as classification rules, PSRs can be used as general behavior rules. In this case two templates are associated. One of them, the antecedent, is used to recognize a set of conditions (a context) which indicates that a particular set of actions is appropriate; when the antecedent template is matched by some event in the environment, the rule is invoked. The second template, the consequent, specifies what actions are to be performed. When the two templates share common parameters, each parameter in the consequent is bound to the same value as the corresponding parameter in the antecedent. These behavior rules may act, for example, as Post productions, transformational grammar rules, or the problem solving rules of STRIPS [7]. In short, a rule with the antecedent A(X) and the consequent C(X) over the variables in the set X is interpreted to mean $(X) [A(X) \Rightarrow C(X)]$. In actual applications, A defines a precondition which can be true of the contents of some working memory, and C defines what is to be done if the precondition is satisfied. Note that any such production can be described by a PSR in which each case relation in the antecedent includes a term of the sort EVENT:a, each case relation in the consequent includes a term of the sort EVENT:c, and the PSR itself includes a case relation {ANTECEDENT:a, CONSEQUENT:c}.

The abstraction of the first and second examples in the sample concept formation task can be represented in the following way:

```

E1*E2:
{ {ABOVE:1,BELOW:2},
  {SAME!SIZE:2,SAME!SIZE:1},
  {SMALL:2},
  {SQUARE:1},
  {SMALL:1},
  {CIRCLE:2},
  {TRIANGLE:3},
  {LARGE:3}}

```

Exemplar 1 is in fact an instance of this abstraction if the parameter 1 is replaced by the parameter b, the parameter 2 by c, and the parameter 3 by a. Likewise, exemplar 2 can be seen to match the abstraction if the parameter 1 is replaced by d, the parameter 2 by f, and the parameter 3 by e.

The comparison problem can be solved by using a technique called interference matching or IM [11-12, 15]. It is a process for identifying all of the common properties of two PSRs and extracting a third PSR which is a template matched by the two exemplars. When two events have N attributes in common, their descriptions will contain at most N case relations which are identical (except for alphabetic differences between the names of corresponding parameters). Figure 2 schematizes IM as a process for finding the intersection containing these case relations.




Figure 2. 'Interference matching.

The circular areas labelled A and B correspond to two PSRs; all of the case relations common to the two PSRs are in the area labelled A*B (read "A star B"). Because any subset of this (conjunctive) set of common relations also defines an abstraction of A and B, it is important to be able to distinguish between the set and its proper subsets. We call any abstraction of A and B which is properly contained in no other abstraction of A and B a maximal abstraction. More formally, if $S(*)A$ denotes that A is a PSR matched by the PSR S, then a maximal abstraction, A, of two PSRs, S and T, satisfies $S(*)A$ and $T(*)A$ and $(B) [B(*)A \ S(*)B \ T(*)B \Rightarrow A(*)B]$.

It should be pointed out that for any two PSRs, there may be more than one abstraction which is maximal in the above sense. For example, given the following two exemplars,

```

E3:    {{CIRCLE:a},
        {RED:a}, {LARGE:a}}

E4:    {{CIRCLE:b}, {CIRCLE:c}, {RED:b},
        {GREEN:c}, {SMALL:b}, {LARGE:c}}

```

two maximal abstractions exist. If the parameters a and b are considered to be identical, the maximal abstraction is

```

E3*E4: {{CIRCLE:1}, {RED:1}}

```

If on the other hand, the parameters a and c are considered to be identical, the maximal abstraction is

```

E3*E4: {{CIRCLE:1}, {LARGE:1}}

```

To perform interference matching on reasonably complex representations, we need an algorithm which, operating within as small a search space as possible, can

discover the best maximal abstractions as quickly as possible. Two approaches to interference matching are known: (1) In the bind-first approach, each parameter in one PSR is associated with a parameter in the second PSR and then a maximal abstraction is found by extracting the case relations which are identical in the two PSRs (modulo the parameter bindings). In this case, if the lesser number of parameters (in either PSR) is MP and the greater number is NP, the number of possible binding functions is combinatorial, (binomial coefficient of NP over MP) * MP!. (2) Alternatively, in the match-first approach, all instantiations of case frames of one type in one PSR are compared with all instantiations of the same type of case frame in the other PSR, and possible parameter bindings are identified by determining which parameters have corresponding properties in comparable relations. Here if NI and MI are the numbers of case relations in the larger and smaller PSR (assuming only one type of case frame), the number of possible ways in which the relations can be forced into correspondence is similarly combinatorial. While it is true that if one were interested in computing abstractions of quite low-level event descriptions (such as undirected graphs) neither method would be much preferable to the other, in most real problems the number of instances of any particular case frame is quite small relative to the number of parameters in the PSR, and so the second method is usually preferable to the first. It is this method which is used in our current work.

The actual algorithm we use has the following form: A randomly selected case relation from one of the exemplar PSRs is put into correspondence with a case relation (which is a parameterization of the same case frame) from a second exemplar PSR; parameters having identical properties are identified as equivalent and the resulting common case relation becomes the (primitive) abstraction associated with that set of parameter bindings. Then other pairs of primitive case relations, one from each of the two exemplar PSRs, are put into correspondence. If a compared pair of relations entails parameter bindings consistent with those already identified, the common relation is added to the abstraction being produced. This new abstraction is the set union of the old abstraction and the new case relation, and the new set of parameter bindings is the set union of those bindings entailed by the previous abstraction and the forced bindings of the parameters in the compared pair of case relations. If a pair of case relations entails parameter bindings inconsistent with those already identified, the common case relation becomes a new (primitive) abstraction.

Clearly, this algorithm may find a number of competing maximal abstractions. Our approach is to build as many distinct abstractions as possible, one relation at a time, until a limitation on the number of distinct abstractions which can be considered at one time is exceeded. At that point, only those abstractions which are most significant in terms of the number and type of case relations they include are retained. These abstractions continue to be extended as other pairs of consistent relations are found; at the same time, the least significant abstractions are continually pruned from further consideration in order to keep the search space as small as possible.

The result of the process is a set of best maximal abstractions, represented as PSRs. Any one of these abstractions (interpreted as existentially quantified) can then be input to SPROUTER together with a third exemplar to produce a set of maximal abstractions of three exemplars, or the process may be repeated on as many additional exemplars as desired. Since a maximal abstraction is compared to an exemplar in the same way that an exemplar is compared to another exemplar, we find

it desirable to store abstractions as PSRs, with the interpretation that their parameters represent existentially quantified variables derived from the correspondence of case relations in the exemplars from which the PSR was induced.

The successive steps involved in producing the maximal abstraction of the first two examples in the concept formation task are shown below.

- (1) {SMALL:1}
- (2) ({SMALL:1}, {ABOVE:2, BELOW:1})
- (3) ((({SMALL:1}, {ABOVE:2, BELOW:1}), {SAME!SIZE:1, SAME!SIZE:2})
- (4) (((({SMALL:1}, {ABOVE:2, BELOW:1}), {SAME!SIZE:1, SAME!SIZE:2}), {SMALL:2})
- (5) (((({SMALL:1}, {ABOVE:2, BELOW:1}), {SAME!SIZE:1, SAME!SIZE:2}), {SMALL:2}), {SQUARE:2})
- (6) (((({SMALL:1}, {ABOVE:2, BELOW:1}), {SAME!SIZE:1, SAME!SIZE:2}), {SMALL:2}), {SQUARE:2}), {CIRCLE:1})

The case relation {SMALL:c} is selected at random from E1 and is then put into correspondence with the case relation {SMALL:f} from E2. The parameters c and f are identified as equivalent and so (since c and f are the first pair of parameters bound) the primitive abstraction {{SMALL:1}} is generated. Then the pair of case relations {ABOVE:b, BELOW:c} and {ABOVE:d, BELOW:f} are put into correspondence. Since the identification of c with f and of b with d is consistent with the already established binding, the primitive abstraction {{ABOVE:2, BELOW:1}} is added to {{SMALL:1}}. It should be noted that our basic IM algorithm actually finds only six of the eight case relations constituting the abstraction. This is because the partial abstraction {{TRIANGLE:3}, {LARGE:3}} was pruned from consideration early in the match under the space limitation constraint. To insure that such complementary relations are not missed, our algorithm, after completing the process described above, searches for additional relations which can extend the abstractions produced. Any such relations which are found are conjoined to the abstraction to produce a maximal abstraction.

SPROUTER, the program which induces abstractions from structural descriptions, is only one part of a classification and learning system which we are developing. The top-level program, called SLIM [4, 10, 16], is a general space limited interference matching procedure which builds abstractions from examples and then uses these abstractions to classify test stimuli. While the abstraction of feature-value representations can be performed by simple bit vector operations (which SLIM itself is capable of), the generation of abstractions from PSRs requires the matching and parameter binding determinations discussed above. The program, SPROUTER, was created for this purpose. Once an abstraction is computed from some PSRs, it is nearly as complex a problem to use it for classification as it was to generate it originally [1, 11-13]. With this in mind, SPROUTER was designed to produce two outputs: one of these is a PSR, which as we have indicated can be matched with subsequent exemplars to produce more refined abstractions; the other is a special purpose recognition network used to exploit an abstraction as a template.

The templates which SPROUTER generates for SLIM are automatically compilable recognition networks or ACORNs [13, 18]. An ACORN is a special data structure, equivalent in representational power to a PSR, but better adapted to serve as a template; it is essentially a Pandemonium pattern recognition system [27], generalized to handle patterns and data described as general

propositional formulae. Once an ACORN has been produced, SLIM can determine whether a descriptive PSR matches it by using the PSR to create an instance list at each of the lowest-level nodes in the ACORN and then allowing the relevant instances of subpatterns of interest to percolate upward in the network. If any instances of the highest-level node are found, the template is matched by the stimulus pattern. The lowest-level nodes of an ACORN correspond to the distinct case frames in a universally quantified PSR and are like the feature demons of a Pandemonium system. A feature demon, however, reports only the number of instances of its particular feature to higher-level demons, whereas the node in an ACORN actually passes its instances up to the higher-level nodes which it supports. The higher-level nodes look for instances of the particular conjunction of case relations in which they are interested, just as higher-level "cognitive demons" in Pandemonium look for specific combinations of feature values. The highest-level node in an ACORN is instantiated if and only if the abstraction is matched by the PSR. Thus this highest-level node corresponds to a Pandemonium's highest-level cognitive demon which recognizes when a pattern of interest is matched. Because ACORNs have been developed to provide a means for sharing the results of the evaluation of subexpressions common to numerous templates, each conjunction of predicates or subtemplates is associated with a single binary-branching node whose two descendants represent the conjoined propositional formulae.

Once a set of best maximal abstractions is computed for two or more exemplars, all training exemplars (or a sample of them) may be examined to see if they match the inferred hypothetical concept or rule. Only to the extent that exemplars of the same class match an abstraction and those of the other classes do not, do we find support for the inference that the abstraction is the criterial concept underlying the training data [9-10]. ACORNs greatly facilitate this examination process. One simply instantiates the terminal nodes of the ACORN whose highest nodes represent the abstractions of interest, and then iteratively computes all instances of each higher-level node from those pairs of instances of its subordinate nodes which satisfy criterial tests on their values. If any instances of the abstraction are produced, the training exemplar matches the abstraction. Without ACORNs, it would be extremely difficult to determine which positive and negative training exemplars matched each abstraction.

III. THE INTERFERENCE MATCHING ALGORITHM

SPROUTER's function, as we have said, is to build ACORNs which can be used by SLIM for recognition. Before this construction process can begin, a set of primitive (bottom-level) nodes must be generated and then instantiated. To generate these nodes, SPROUTER reads in the set of case frames which are relevant to the task it is facing. For each of these case frames, a primitive node is created which is essentially a universally quantified case relation. SPROUTER then finds, in the descriptive PSRs of two exemplars, the set of distinct instances (case relations) which are instances of each of these nodes. Each node has two associated instance lists; each of these lists contains the instances of the case relation for one of the exemplars. For example, given the two case frames N1: {CIRCLE}, N2: {ABOVE, BELOW} and the two exemplars

E5: {{CIRCLE:a, CIRCLE:b},
{ABOVE:a, BELOW:b}}

E6: {{CIRCLE:c}}

SPROUTER will create two nodes, N1 and N2, and then

produce four instance lists. Two of these lists, ($E5/a$), ($E5/b$) and ($E6/c$), are associated with node N1. The other two, ($E5/a$, $E5/b$) and (), are associated with node N2.

When the primitive nodes have been instantiated, SPROUTER produces the set of maximal abstractions of the two PSRs by constructing, bottom-up, a binary-branching ACORN. Each higher-level node of this network is a conjunction of two nodes, one of which is always a primitive node. Before initiating the building process, SPROUTER deletes all of the primitive nodes that do not have at least one instance from each exemplar. Then one exemplar, the one with fewer instances over the remaining nodes, is tagged E_{intro} ; the other exemplar is tagged E_{comp} . And each instance of E_{intro} is marked as unused. SPROUTER then begins the actual construction. An unused E_{intro} instance from a primitive node is chosen as one of the two instances to be used in the construction; it is selected on the basis of the likelihood of its being an instance of a node which is a constituent of a best maximal abstraction. This instance is then paired with every instance from E_{intro} of every node. Each of these pairs of instances is used to construct a candidate node which will accept instance pairs only if they are equivalent to the prototypic pair. If there is at least one such pair of instances in E_{comp} , the candidate node is added to the network and all instances of the node (from both exemplars) are computed. Thus, each step in the abstraction building process involves combining, iteratively, an unused instance from a primitive node with each other instance in the ACORN. After each of the resulting conjunctive nodes is generated for a pair of instances from E_{intro} , all instances of that node, first from E_{comp} and then from E_{intro} , are computed. If no instances are found in E_{comp} , the node represents an abstraction that is not true of the second exemplar and so the node is not added to the network. The process continues until all of the case relations that are common to both exemplars have been conjoined.

Of course, this algorithm, left unconstrained, would build a node for each subset of case relations in E_{intro} for which there was an equivalent subset in E_{comp} . Clearly, the size of the search space would increase exponentially. Thus, for even small problems, it is important to somehow reduce the number of nodes constructed. We use two heuristics. The first of these enables us to keep the search space to a manageable size by providing for the automatic pruning of those conjunctions least likely to be part of a best maximal abstraction. To determine which partial abstractions are least promising, a value is computed which we call the utility of a node. Basically, the utility of a node is an increasing function of the number of properties covered by the node and a decreasing function of the number of distinct parameters needed to instantiate the node. More specifically, our current utility measure adds 1.0 for each property of a case relation and subtracts 1.0 for each distinct parameter in the associated PSR. Our justification for this rather rough measure of utility is that it will yield as the highest valued nodes, those with the greatest scope and connectivity. Equivalently, the higher the utility of a node, the more informative and apparently "better" it is as an abstraction.

During the construction of the ACORN, a list of all nodes currently in the network is maintained. This list, which is ordered by the utility of its elements, has a stipulated maximum length. Whenever the number of total nodes in the ACORN exceeds this stipulated maximum, a primitive node which does not support any higher-order nodes is marked as removed from consideration. If all remaining primitive nodes support some higher-level node, then the least valued maximal abstraction (provided there

is more than one maximal abstraction in the network) and all nodes supporting it (or supporting one of its supports, recursively) and not supporting some other higher valued maximal abstraction are deleted (or marked as removed from consideration if they are primitive nodes). Thus, the number of nodes in the network can exceed the stipulated maximum only if just one maximal abstraction remains. While in some cases, it might be desirable to require that at least k ($k > 1$) best maximal abstractions be maintained, we have not yet found a need for this option.

As a result of the limitation on nodes in the ACORN, the typical behavior during construction is as follows: Instances are introduced one-at-a-time from E_{intro} and are conjoined with other E_{intro} node instances to form PSRs representing subsets of case relations of varying utility. As soon as the number of nodes corresponding to these nodes in the ACORN exceeds the stipulated maximum, the maximal node with the lowest utility together with all nodes which support only it are deleted from the network. This construction-and-pruning cycle is repeated until the set of best maximal abstractions has been found.

The second heuristic provides the search with direction by indicating which one of the unused instances is to be used in the next cycle of construction. Our search for the best maximal abstractions is essentially hill climbing, but occurs on many hills simultaneously. Since our pruning heuristic enables us to maintain a gradually decreasing number of maximal abstractions, the number of hills under consideration is reduced as the search progresses. Clearly, if we could select first all of those instances from E_{intro} which were instances of the best maximal abstractions (the highest hills), then our search, since it would take place in an essentially unimodal space, would be as efficient as possible. Of course it is impossible to determine a priori which instances are instances of the best maximal abstractions. However, by using a variant of the utility function described above, it is possible to compute, fairly cheaply, the upper bound of the actual utility of any node which might be constructed. Using this strategy, we can, at relatively little cost, significantly increase the probability that the node constructed will be a constituent of a best maximal abstraction. The selection procedure we use is as follows: We set a sampling factor (currently 20%) for the proportion of the unused instances from E_{intro} which are to be examined. We select at random this percent of the unused instances (but at least three until there are fewer than three unused instances). For each of the instances in this sample, we determine an upper bound of the utility of all of the nodes which could be constructed by conjoining the sampled instance with the remaining instances of nodes still under consideration. The one instance which produces the node with the highest potential utility is constructed.

The actual construction of a node is a two step process. First SPROUTER creates a set of tests which are both necessary and sufficient to accept just those instances that are equivalent to the pair of instances used as a model in building the higher-level candidate node. It is possible to create such a set of tests working only with the sameness or difference of selected parameters. For example, to construct an ACORN node to accept the two instances {CIRCLE:c} and {ABOVE:a, BELOW:c}, a same parameter (SP) test is generated to insure that the first parameter of the first case relation is the same as the second parameter of the second relation, and a different parameter (DP) test is generated to insure that no non-explicit SPs are accepted. If one thinks of this ACORN node as being constructed from a left and a right instance, where the parameter of the left instance is numbered 1, and the parameters of the right instance are numbered 2 and 3, then a minimal complete set of tests needed to

exactly represent the same and different relations are {SP:1, SP:3} and {DP:1, DP:2}.

After the set of tests has been created, the candidate node is associated with a generator set which specifies how the parameters of its instances are to be extracted from pairs of subordinate instances which satisfy the node's SP and DP tests. Because of the implicit requirement for DP relations to hold on all distinct parameters, the order of the new relation is exactly the number of distinct parameters in the two relation instances used in building the node. In the above example, there would be two parameters in each instance of the new node and these would correspond to parameters 1 and 2 (since 1 and 3 are identical). The generator list for this node would be just (1,2). From the nature of the explicit SP and DP tests used, it follows that any two nodes having instances derived from equivalent pairs of instances must be equivalent. Whenever such a duplicate node is constructed, it is removed from the ACORN.

It should be apparent that an ACORN constructed in the fashion described above will not necessarily contain a maximal abstraction. Whether or not it will is partially dependent on what maximum has been stipulated for the number of nodes in the ACORN. But even if the stipulated maximum is large enough so that the highest node in the ACORN is a constituent of a maximal abstraction, the ACORN may not be complete; that is, some of the case relations in the abstraction may have been lost. This can occur if one or more primitive nodes whose instances are a part of the abstraction were removed from consideration early in the construction process. In such a case, however, it is always possible to extend the ACORN with conjunctions of these lost primitive node instances. This is done by successively reintroducing into the construct-and-prune cycle each instance in E_{intro} which does not support all of the instances of all of the highest nodes in the ACORN. Each re-introduced instance is conjoined with each of the instances of each highest node to produce candidate nodes. If instances of any of these new abstractions are found in E_{comp} , these new nodes are retained; the ACORN is then extended further, in the same way, until the best maximal abstractions have been found.

We have already seen the abstraction SPROUTER constructs given the first two exemplars in the first concept formation task. The set of case frames from which the primitive nodes were created, all three exemplars, and the best maximal abstraction found by SPROUTER are given below.

```
CF:
{ N1:{CIRCLE},
  N2:{SQUARE},
  N3:{TRIANGLE},
  N4:{LARGE},
  N5:{SMALL},
  N6:{INNER, OUTER},
  N7:{ABOVE, BELOW},
  N8:{LEFT, RIGHT},
  N9:{SAME!SHAPE, SAME!SHAPE},
  N10:{SAME!SIZE, SAME!SIZE},
  N11:{BESIDE, BESIDE},
  N12:{CONTIGUOUS, CONTIGUOUS}}
```

```
E1:
{ {TRIANGLE:a, SQUARE:b, CIRCLE:c},
  {LARGE:a, SMALL:b, SMALL:c},
  {INNER:b, OUTER:a},
  {ABOVE:a, ABOVE:b, BELOW:c},
  {SAME!SIZE:b, SAME!SIZE:c}}
```

```
E2:
{ {SQUARE:d, TRIANGLE:e, CIRCLE:f},
  {SMALL:d, LARGE:e, SMALL:f},
  {INNER:f, OUTER:e},
  {ABOVE:d, BELOW:e, BELOW:f},
  {SAME!SIZE:d, SAME!SIZE:f}}
```

```
E3:
{ {SQUARE:g, CIRCLE:h, CIRCLE:i},
  {SMALL:g, LARGE:h, SMALL:i},
  {INNER:i, OUTER:h},
  {ABOVE:g, BELOW:h, BELOW:i},
  {SAME!SHAPE:h, SAME!SHAPE:i},
  {SAME!SIZE:g, SAME!SIZE:i}}
```

```
E1+E2+E3:
{ {N10:{SAME!SIZE:1, SAME!SIZE:2},
  {N7:{ABOVE:1, BELOW:2},
  {N1:{CIRCLE:2},
  {N5:{SMALL:1},
  {N5:{SMALL:2},
  {N2:{SQUARE:1},
  {N4:{LARGE:3}}
INSTANCES FROM EXEMPLAR E1+E2
([E1+E2/2, E1+E2/1, E1+E2/3])
INSTANCES FROM EXEMPLAR E3
([E3/g, E3/i, E3/h])
```

SPROUTER took 6 seconds of cpu time on a PDP-10 (model KA-10) to produce $E1+E2$ which it found after constructing 14 nodes (7 more than necessary). SPROUTER took 3 seconds and constructed 6 nodes (the fewest possible) to produce $(E1+E2)+E3$. The abstraction SPROUTER found, however, though it is the best abstraction producible using our match-first method, is not maximal. It is missing two case relations. As we indicated in the first section of the paper, the abstraction SPROUTER induces is the following:

There are three objects, including a small circle and a small square. The square is above the circle. The third object is large.

The best maximal abstraction includes the specification that the large object contains another one which is one of the two small objects. SPROUTER is unable to find this abstraction for two reasons: (1) The grain size of the representations used in describing the examples is too big; more atomic uniform representations are needed to make abstraction, which is a strictly subtractive process, more generally applicable [11, 15]. (2) Many-one parameter correspondences must be allowed in order to insure that relevant correspondences are not lost. These two problems, whose solution requires methods of greater generality than we have currently implemented, are discussed in detail in the CMU technical report from which this paper is taken.

IV. CONCLUDING REMARKS

SPROUTER has already solved learning problems of theoretical significance and of considerable complexity. Because of the extensive size of the search spaces, such learning could not be done with simple enumerative matching algorithms. In essence, SPROUTER establishes the feasibility of induction from non-trivial exemplar descriptions. In many respects, however, SPROUTER is quite primitive. It is a purely syntactic matcher; it knows nothing at all about the underlying structure or significance of any of the predicate descriptions upon which it operates. For this reason, its utility function, and thus its heuristics, are very weak. One interesting approach to improving the performance of SPROUTER would be to

provide it with domain-specific utility functions. For example, if SPROUTER knew that concordance on antecedent or consequent relations was more important than concordance on most other relations, it would never attempt to match the antecedent part of an example with a consequent part. Similarly, if it knew that concordance of higher-order grammatical constructs (e.g., a sentence) was more significant than concordance on lower-order ones, it could quickly zero in on the concordances of two sentence structures and then continue building abstractions in an essentially top-down fashion.

Even though SPROUTER's performance has been quite impressive on several tasks, there are three difficulties impeding the use of such a learning machine in general applications. First, an empirical question has been raised regarding the preferability of approaches to induction based on the one-one and many-one binding alternatives. If object integrity in representations is generally tenuous -- that is, if each object in one PSR can correspond to multiple, diverse objects in another PSR, as was the case in the concept formation task described above -- abstraction procedures based on the many-one approach will have to be developed. Second, one must identify which real-world problems can be solved by interference matching methods. Because the case frames which SPROUTER uses in inferring abstractions are assumed to be externally provided, the utility of our basic method depends upon the prior identification of useful properties. Third, while SPROUTER uses what is basically a common heuristic search strategy in AI (breadth-first with pruning of low-valued paths), it is still too expensive to be widely applied in practical problems.

REFERENCES

- Barrow, H. G., Ambler, A. P., and Burstall, R. M. Some techniques for recognizing structures in pictures. In *Frontiers of Pattern Recognition*, S. Watanabe (Ed.), Academic Press, New York, 1972.
- Bobrow, D., and Winograd, T. A knowledge representation language. *Cognitive Science*, in press.
- Bruner, J. S., Goodnow, J. J., and Austin, G. A. *A Study of Thinking*. Wiley, New York, 1956.
- Burge, J., and Hayes-Roth, F. A novel pattern learning and classification procedure applied to the learning of vowels. *Proc. of the 1976 IEEE Intl. Conf. Acoustics, Speech, and Signal Processing*. Philadelphia, 1976.
- Ernst, G. W., and Newell, A. *GPS: A Case Study in Generality and Problem Solving*. Academic Press, New York, 1969.
- Evans, T. G. A program for the solution of geometric-analogy test questions. In M. Minsky (Ed.), *Semantic Information Processing*. MIT Press, Cambridge, 1968.
- Fikes, R. E., and Nilsson, N. J. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2, (1971), 189-208.
- Gallot, F. *Inquiries into Human Faculty and its Development*. Dent, London, 1907.
- Hayes-Roth, F. A structural approach to pattern learning and the acquisition of classificatory power. *Proc. First Intl. Jt. Conf. Pattern Recognition*, 1973.
- Hayes-Roth, F. Schematic classification problems and their solution. *Pattern Recognition* 6, 2 (Oct. 1974), 105-114.
- Hayes-Roth, F. Fundamental mechanisms of intelligent behavior: the representation, organization, acquisition, and use of structural knowledge in perception and cognition. Doctoral Dissertation, The University of Michigan, Ann Arbor, 1974.
- Hayes-Roth, F. An optimal network representation and other mechanisms for the recognition of structured events. *Proc. Second Intl. Jt. Conf. Pattern Recognition*, 1974.
- Hayes-Roth, F. Representation of structured events and efficient procedures for their recognition. *Pattern Recognition*, 8, 3 (Aug. 1976), 141-150.
- Hayes-Roth, F. Patterns of induction and associated knowledge acquisition algorithms. In *Pattern Recognition and Artificial Intelligence*, C. Chen (Ed.), Academic Press, New York, in press.
- Hayes-Roth, F. Uniform representations of structured patterns and an algorithm for the induction of contingency-response rules. *Information and Control*, 32, (1976), in press.
- Hayes-Roth, F., and Burge, J. Characterizing syllables as sequences of machine-generated labelled segments of connected speech: a study in symbolic pattern learning using a conjunctive feature learning and classification system. *Proc. Third Intl. Jt. Conf. Pattern Recognition*, 1976.
- Hayes-Roth, F., Erman, L. D., Fox, M., and Mostow, D. J. Syntactic processing in Hearsay-II. In *Speech understanding systems: summary of results of the five-year research effort*. Department of Computer Science, Carnegie Mellon University, Pittsburgh, 1976.
- Hayes-Roth, F., and Mostow, D. J. An automatically compilable recognition network for structured patterns. *Proc. Fourth Intl. Jt. Conf. Artificial Intelligence*, 1975.
- Hirschman, L., Grishman, R., and Sager, N. Gramatically based automatic word class formation. *Information Processing and Management*, 11, 1-2 (Mar. 1975), 39-57.
- Hunt, E. B. *Concept Formation: An Information Processing Problem*. Wiley, New York, 1962.
- Michalski, R. S. AQVAL/I -- Computer implementation of a variable valued logic system VLI and examples of its application to pattern recognition. *Proc. First Intl. Jt. Conf. Pattern Recognition*, 1973.
- Minsky, M. A framework for representing knowledge. In *The Psychology of Computer Vision*, P. H. Winston (Ed.), McGraw-Hill, New York, 1975.
- Moore, J., and Newell, A. How can Merlin understand? In *Knowledge and Cognition*, L. W. Gregg (Ed.), Erlbaum, New York, 1974.
- Plotkin, G. D. A note on inductive generalization. In *Machine Intelligence*, vol. 5, B. Meltzer and D. Michie (Eds.), American Elsevier, New York, 1970.
- Plotkin, G. D. A further note on inductive generalization. In *Machine Intelligence*, vol. 6, B. Meltzer and D. Michie (Eds.), American Elsevier, New York, 1971.
- Reed, S. K., Ernst, G. W., and Banerji, R. The role of analogy in transfer between problem states. *Cognitive Psychology* 6, 3 (Nov. 1974), 436-450.
- Selfridge, O. G. Pandemonium: a paradigm for learning. Symposium on the Mechanisation of Thought Processes. H. M. Stationery Office, 1959.
- Shaw, A. C. Picture graphs, grammars, and parsing. In *Frontiers of Pattern Recognition*, S. Watanabe (Ed.), Academic Press, New York, 1972.
- Stoffel, J. C. A classifier design technique for discrete variable pattern recognition problems. *IEEE Trans. on Computers* C-23, 4 (Apr. 1974), 428-441.
- Treliakoff, A. Computer-generated word classes and sentence structures. *Information Processing 1974, Proc. IFIP Congress 74*, 1974.
- Vere, S. A. Induction of concepts in the predicate calculus. *Proc. Fourth Intl. Jt. Conf. Artificial Intelligence*, 1975.
- Winston, P. H. Learning structural descriptions from examples. In *The Psychology of Computer Vision*, P. H. Winston (Ed.), McGraw-Hill, New York, 1975.