

DISCOVERING PATTERNS IN  
SEQUENCE OF EVENTS

by

*Tom Dietterich*  
*Ryszard S. Michalski*

Artificial Intelligence Journal, Volume 25, No. 2, 1985.

Volume 25 Number 2 February 1985

ISSN 0004-3702

AINTBB 25(2) (1985) 117-246

# Artificial Intelligence

AN INTERNATIONAL JOURNAL



North-Holland  
Amsterdam

# Discovering Patterns in Sequences of Events

---

**Thomas G. Dietterich**

*Department of Computer Science, Stanford University, Stanford,  
CA 94305, U.S.A.*

**Ryszard S. Michalski**

*University of Illinois, Urbana, IL 61801, U.S.A.*

Recommended by Saul Amarel

---

## ABSTRACT

*Given a sequence of events (or objects), each characterized by a set of attributes, the problem considered is to discover a rule characterizing the sequence and able to predict a plausible sequence continuation. The rule, called a sequence-generating rule, is nondeterministic in the sense that it does not necessarily tell exactly which event must appear next in the sequence, but rather, defines a set of plausible next events.*

*The basic assumption of the methodology presented here is that the next event depends solely on the attributes of the previous events in the sequence. These attributes are either initially given or can be derived from the initial ones through a chain of inferences. Three basic rule models are employed to guide the search for a sequence-generating rule: decomposition, periodic, and disjunctive normal form (DNF). The search process involves simultaneously transforming the initial sequences to derived sequences and instantiating models to find the best match between the instantiated model and the derived sequence. A program, called SPARCIE, is described that implements most of the methodology as applied to discovering sequence generating rules in the card game Eleusis. This game, which models the process of scientific discovery, is used as a source of examples for illustrating the performance of SPARCIE.*

---

## 1. Introduction

Inductive learning—that is, learning by generalizing specific facts or observations—is a fundamental strategy by which we acquire knowledge about the world. This form of learning is rapidly becoming one of the central research topics in AI. Most research on computer models of inductive learning has addressed the problem of inducing a general description of a concept from a

*Artificial Intelligence* 25 (1985) 187–232

collection of independent instances of that concept (the so-called training instances). Thus, the research has dealt with learning concepts that represent a certain class of instances. The instances can be specific physical objects, processes, actions, images and so on. The learned concepts are general descriptions of given classes of such instances. These descriptions provide a means for assigning an unknown instance to the appropriate class (or concept).

Learning problems of this type include:

- learning a checkers evaluation function [29, 30] that assigns to a given class of board situations a certain value;
- learning descriptions of block structures [36];
- determining rules for interpreting mass spectrograms [5];
- formulating diagnostic rules for soybean diseases [25]; and
- discovering heuristics to guide the application of symbolic integration operators [27].

In Samuel's checkers program, for example, each training instance was a board situation represented as a vector of 16 attributes. The learned concept was an evaluation function that computed the 'value' of any board position for the side whose turn it was to move. No relationships between different board positions were considered. Similarly, Michalski's AQ11 program was given independent training instances, each describing a diseased soybean plant in terms of 35 multi-valued attributes. Each plant had one of 19 common soybean diseases. From several hundred training instances, the program inferred general diagnostic rules for these diseases. These rules outperformed rules developed through consultation with expert plant pathologists [25].

In all of the above problems, the system learns a general class description from instances of the class, and therefore, this type of inductive learning can be called *instance-to-class generalization*. A review of several methods for such instance-to-class generalization can be found in [8, 24].

Another type of inductive learning involves constructing a description of a whole object by observing only selected parts of it. For example, given a set of fragments of a scene, the problem is to hypothesize the description of the whole scene. A very important case of such *part-to-whole generalization* is where the 'part' consists of a fragment of a sequence of objects (or a process evolving in time) and the problem is to induce the hypothetical description of the whole sequence (the process). Once such a description is found, it can be used to predict the possible continuations of the given sequence or process. This class of part-to-whole inductive learning problems we will call *prediction problems*.

Examples of such prediction problems include:

- predicting the outcomes of medical treatments;
- predicting the behavior of stock and futures markets; and
- estimating the damage to corn from cutworm infestation [4].

This paper investigates the prediction problem for a sequence of events (or

objects) that are characterized by a finite set of discrete-valued attributes. An elementary problem of this type is letter-sequence prediction, in which each object in the sequence is characterized by only one attribute: the name of the letter. For example, given a sequence of letters such as

A B X B C W C D V . . .

the learning program must discover a 'pattern'—that is, a rule that governs the generation of letters in the sequence. In this case, such a rule might state that the sequence is a periodically repeating subsequence of three letters in which the first letter is the successor of the corresponding letter in the previous 3-letter subsequence, the second letter is the successor of the first, and the third letter is the predecessor of the corresponding letter in the previous 3-letter subsequence. Early papers by Simon and Kotovsky [15, 32, 33] show that just a few relationships (such as successor, predecessor, and equality) are sufficient to represent most such patterns. Related work by Solomonoff [34] and Hedrick [13] has investigated grammatical approaches to describing letter sequences.

The sequence prediction problem becomes more difficult when the sequence consists not of simple objects with only a single relevant attribute (like the problem just described), but instead of objects with many relevant attributes. Further complexity is introduced if the pattern describing the sequence also involves a variety of relationships among these attributes. For example, the pattern may involve the periodicity of recurrence of certain properties or the dependence of the next object in the sequence on the properties of objects preceding it at some arbitrary distance in the past. A sequence prediction problem exhibiting the above-mentioned complexities arises in the card game Eleusis, which models the process of scientific discovery [1, 12]. We will use examples from this game to illustrate our general methodology for discovering patterns in event sequences. The rules for Eleusis are briefly explained in Section 2.1.

The fundamental assumption underlying the methodology is that the 'correct' sequence-generating rule involves only attributes of the events in the sequence. Those attributes can be either initially given to the program or derived by the program from the initial ones through a sequence of inferences. If this assumption does not hold, the method can still produce a sequence-generating rule, but with little predictive power.

Before we discuss further the problem of discovering sequence-generating rules, let us first show its significance to current AI research. Its significance stems from the fact that its solution involves addressing three major AI problems:

- (1) the representation problem;
- (2) the problem of performing *model-driven inductive learning* with multiple models; and
- (3) the problem of reasoning about temporal processes.

The specific representation problem of interest here is that of automatically determining an appropriate series of transformations of the initial sequence description so that the pattern can be found. The multiple-model inductive learning problem arises because no single model can provide sufficient guidance to the search for plausible descriptions in this domain. The relationship of this problem to reasoning about temporal processes is not as strong as the other two problems. However, since such processes include as a special case discrete-time linear sequences, some of the techniques developed for event sequence prediction should be relevant to reasoning about continuous temporal processes.

In the next two sections, we discuss in detail the representation problem and the problem of multiple-model induction as they arise in this domain.

### 1.1. Task-oriented transformation of description space

The problem of transforming the initial problem description arises in many practical domains in which the given data (e.g., the training instances in concept learning) are observations or measurements that do not include the information directly relevant to the task at hand. For example, in character recognition, the input typically consists of a matrix of light intensities representing a character, but the relevant information includes position-invariant properties of letters such as the presence of a line on the left or right of a character, occurrence of line endings, closed contours, and so on (e.g., [14]). These position-invariant properties can be made explicit by applying *task-oriented* transformations to the raw data.

An example of a learning program that performs task-oriented transformations is INTSUM (a part of the Meta-DENDRAL system [5]). INTSUM is presented with raw training instances in the form of chemical structures (graphs) and associated mass spectra (represented as fragment masses and their intensities). For each fragment in the mass spectrum, INTSUM must determine the bonds that could have broken to produce that fragment. A simple mass spectrometer simulator is used to develop these hypothesized bond breaks. Each of the resulting transformed training instances has the form of a chemical structure and a set of bonds that broke when that structure was placed in the mass spectrometer. It is this information that is provided to the remaining parts of the Meta-DENDRAL system (programs RULEGEN and RULEMOD).

In character recognition programs and in Meta-DENDRAL, the data transformations are fixed in advance. Future learning systems, however, may not know the proper transformations *a priori*. These learning systems will need to select or invent appropriate task-oriented transformations for each learning situation.

This description-space transformation problem has been called by various authors the *data interpretation problem* [8] or the *reformulation problem* [2]. We

prefer the term *task-oriented transformation problem*, since it emphasizes that the proper choice of data transformations depends upon the task being performed. In the prediction problem discussed here, the desired sequence-generating rules are described in a language quite different from the language used to describe the raw sequence. The learning system searches for appropriate data transformations from four a priori defined classes of transformations. These transformations are applied to the raw sequence to produce a transformed sequence that is more amenable to pattern discovery.

The task-oriented transformation problem is part of a spectrum of problems faced by learning programs. The simplest learning algorithms (e.g., linear regression) determine the coefficients for a predetermined, fixed set of variables combined in a predetermined expression. More sophisticated are learning algorithms, such as the  $A^9$  algorithm [21, 22] and the candidate-elimination algorithm [26, 27], that are able to determine which terms are relevant and how they should be combined. A third level of sophistication is obtained by learning programs that perform interpretative transformations (e.g., [35], Meta-DENDRAL [5]). These programs augment the basic inductive algorithms by applying a set of predetermined transformations to the data prior to inductive generalization. The next step of difficulty is represented by learning algorithms that select description space transformations under the guidance of special heuristics. Very few researchers have addressed this problem (e.g., [17, 23]). The most sophisticated algorithms currently envisioned—but not yet developed—would be capable of discovering new data transformations. Table 1 shows this spectrum of inductive learning problems.

TABLE 1. Spectrum of learning problems in increasing order of difficulty

- 
1. Determine coefficients
  2. Select relevant variables and combine
  3. Apply predetermined transformations
  4. Select transformations under heuristic guidance
  5. Discover new transformations
- 

The method presented in this paper falls under category 4 (i.e., select transformations under heuristic guidance), since it searches four general classes of transformations and employs heuristics reflecting domain-specific knowledge.

## 1.2. Learning with multiple models

The second major problem that arises in this prediction problem is concerned with using multiple description models in the process of inductive learning. This problem has not received much attention in previous AI research. Most

existing systems employ a *single model* that provides guidance to the induction algorithm as it searches a space of possible descriptions. Many systems, for example, use conjunctive models to represent concepts, that is, they assume that the learned concept is representable as a conjunction of predicates. By constraining the search to consider only conjunctive descriptions, the learning problem is greatly simplified. A more general approach, employed by Michalski [21, 22], constrains descriptions to be in disjunctive normal form with *fewest* conjunctive descriptions. This constraint is satisfied (approximately) by having the induction algorithm find first one conjunction, and then another, and so on until all of the training instances are covered. Meta-DENDRAL employs a fairly elaborate simulation model of the operation of the mass spectrometer to guide its search for cleavage rules [5]. In general, all of these systems use a single model, and very few authors have made their models explicit.

One researcher who has employed multiple models is Persson [28]. He applied four different models to the problem of extrapolating number- and letter-sequences. Briefly, these models were

(1) a model that computes the coefficients and the degree of a polynomial by applying Newton's forward-difference formula (the degree can be arbitrarily large);

(2) an extended model that discovers exponential rules of the form  $ABC^C$ , where  $A$  is a polynomial of degree 4 or less and  $B$  and  $C$  are polynomials of degree 1 or less (i.e.,  $B$  and  $C$  are of the form  $ax + b$ );

(3) a simple periodic model for periods of length 2 (i.e., intertwined sequences); and

(4) a generalization of the Kotovsky and Simon model for Thurstone letter-series that can discover simple periodic and segmented sequence-generating laws.

These models are applied in an artificial learning situation in which the program is given a *sequence* of sequence-extrapolation problems. Thus, in addition to attempting to solve each individual sequence-extrapolation problem, Persson's program tries to predict the *kind* of sequence-problem that it will next receive.

Persson's work shows the value of employing multiple description models to search for sequence-generating rules. The major limitation of Persson's approach, however, is that it is specific to number- and letter-sequence prediction. His methods cannot solve the more general prediction problem described in this paper in which events have multiple attributes (numerical and non-numerical) and the sequences are characterized by nondeterministic logic-based sequence-prediction rules.

Table 2 shows a spectrum of five model-based learning methods. The simplest approach is to use a single fixed model. This has been the common approach in AI thus far. The next step is to provide a learning program with a set of models from which it would choose the most appropriate ones. This is



TABLE 2. Spectrum of model-based methods in increasing difficulty

- 
1. Single model
  2. Selection from a few models
  3. Predetermined generation of models
  4. Heuristically-guided generation of models
  5. Discovery of new models
- 

the approach used by Persson. The third level of sophistication would be to have the program generate a predetermined set of models, just as the learning program applies a predetermined set of data transformations. This could be improved by having the program decide which models to generate on the basis of special heuristics. Finally, an even more sophisticated program would be able to invent new models and apply them to guide the learning process.

The approach described in this paper searches a predetermined space of possible models in a depth-first fashion, and hence, falls under point 3 of Table 2.

The main theoretical contributions of this research include the development of techniques for solving the above two problems—that is, selection of task-oriented transformations and application of multiple description models—and also the processes for matching instantiated models to the transformed sequences using a bidirectional search.

The following sections of this paper

- (1) define the sequence prediction problem under consideration;
- (2) describe the methods used for representing and transforming the initial training instances;
- (3) present techniques for representing the models and sequence-generating rules; and finally,
- (4) provide the details of the program SPARC/E, which implements most of the described methodology. The program is illustrated by a few selected examples of its operation when applied to the inductive card game Eleusis.

## 2. Problem Statement

Suppose we are observing a process that generates some objects, one after another, and arranges them into a sequence. Suppose that the objects are generated from a known set and that there exists an underlying law that specifies at least some of the properties of every new generated object. We will call such a law a *sequence-generating rule*. It is assumed that the law is expressed in terms of properties that are either observable properties of objects present in the sequence up to the moment when a new object is generated or properties that can be derived from such observable properties by some known inference rules.

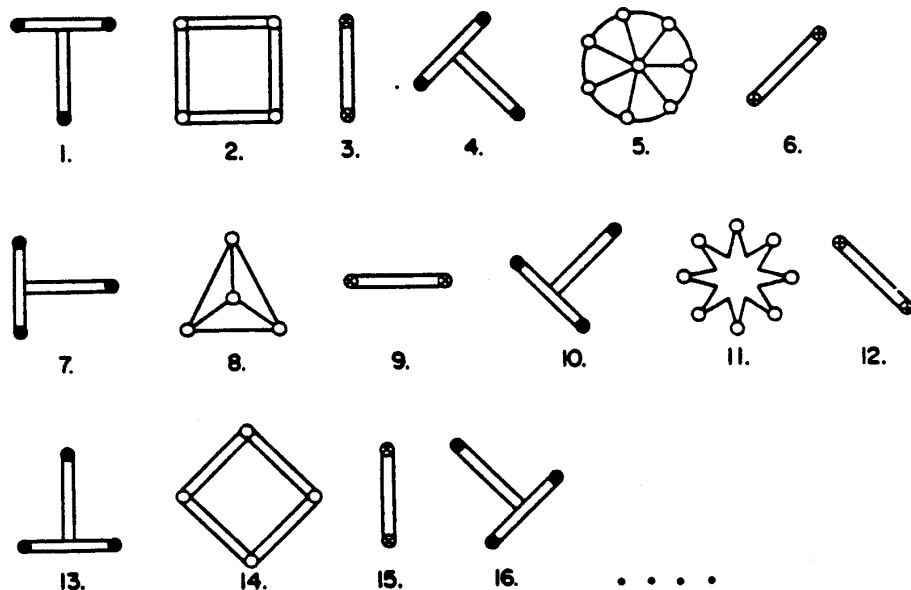
We are interested in the most general kind of sequence-generating law in which the law does not necessarily completely determine which objects can or cannot

appear next in the sequence. The law merely states some properties that constrain the next object to be a member of a restricted set. Thus, such a generating rule is nondeterministic. The task of discovering such a generating law is a difficult learning task, requiring task-specific data transformations and model-guided induction. We will call this learning problem a *nondeterministic prediction problem* (NDP, for short). If the law guiding the generation of the sequence completely defines the next object at every point in the sequence, then the NDP problem reduces to a *deterministic prediction problem* (DP, for short). In the DP problem, it is assumed that there is no randomness in the generation of the next object. The next object is strictly a function of the past objects.

Many researchers have previously considered DP problems such as letter-sequence prediction, number-series extrapolation, economical prediction, and prediction of the behavior of a computer system. All of the methods developed make strong use of the constraint that in a DP problem, *all* attributes of the next object in the sequence are determined by the previous objects in the sequence. The NDP problem is more difficult to solve, because only a partial description of the original sequence is sought. Consequently, many more plausible hypotheses must be considered during the inductive learning process.

Let us illustrate a simple NDP problem by an example. Suppose we are given a snapshot of an ongoing process that has already generated the objects (graphs) shown in Fig. 1.

The observable properties of each graph are: the NUMBER OF NODES, the SHAPE of the graph (T-junction, square, bar, wheel, triangle, star, diamond), the TEXTURE of each node (solid black, blank, and cross), and the ORIENTATION of the graph (applicable only to graphs that are elongated in some direction, expressed as



degrees clockwise from vertical). Suppose we would like now to predict what could be the next object.

By examining the given sequence in Fig. 1, we can observe that it can be partitioned into segments of three graphs in length. The nodes of the graphs in each triplet have TEXTURE in the order (solid black, blank, cross). The SHAPES of the graphs are always (T-junction, \*, bar) (where \* denotes any shape). We can also notice that the ORIENTATION of the T-junction changes by  $-45$  degrees each time, while the ORIENTATION of the bar increases by  $+45$  degrees each time. Finally, the NUMBER OF NODES in the center graph alternates between 4 and 8. If the above regularities indeed constitute the generating law, we can hypothesize that the next graph in the sequence will have 8 blank nodes, and then after that there will be a graph that is a slanted bar with crossed nodes and ORIENTATION of 225 degrees (slant downwards to left). Thus, with regard to the first predicted object, we know only two properties (NUMBER and TEXTURE of nodes), and with regard to the second predicted object, we know it completely. It is easy to see that the problem of letter-sequence prediction is a special case of the NDP problem where each object is a letter of an alphabet with a single attribute, which is its name. It also has one derived attribute that is its position in the alphabet. (The order of letters in the alphabet is externally-provided domain knowledge.) In a typical letter-series prediction problem, it is assumed that the sequence can be continued in only one way. Therefore, it is a DP problem. Since each letter is defined completely by specifying its name (or its position in the alphabet), the sequence-generating rule is a function of the values of just one attribute.

### 2.1. An exemplary NDP problem: the card game Eleusis

An interesting NDP problem occurs in the card game Eleusis, invented by Robert Abbott [1, 12]. Eleusis is an inductive game in which players attempt to discover a generating rule for a sequence of cards. This 'secret rule' is invented by the dealer before the game. Each player, in his or her turn, adds one card to the sequence, and the dealer indicates whether the card is a correct (or incorrect) extension of the sequence (i.e., satisfies or does not satisfy the secret sequence-generating rule). Players who play incorrectly are penalized by having additional cards added to their hands. The goal of each player is to get rid of all of the cards in his hand, which is only possible if 'correct' cards are played. The cards played during the game are displayed in the form of a layout in which the correct cards form a 'main line' and incorrect cards form 'side lines' branching down from the main line at the card that they followed. Fig. 2 shows a typical Eleusis layout for the sequence-generating rule 'Play alternating red and black cards'. In this game, the 3 of hearts was played first, followed by a 9 of spades and a Jack of diamonds. All of these were correct. Following the Jack, a 5 of diamonds was played. It appears on a sideline below the Jack, because it was not a correct extension of the sequence. (At this point a black card is required.) The 4 of clubs was then correctly played, and so

Main line:	3H	9S	JD	4C	JD	2C	10D	2C	5H
Side lines:			5D		AH	AS	8H		
					8H	10S	7H		
					QD		10H		

FIG. 2. A sample Eleusis layout.

Eleusis provides a good domain for studying the use of task-oriented data transformations to aid learning. Frequently, the generating law for an Eleusis sequence is stated in terms of descriptors that are not present in the initial sequence. In this example, for instance, the generating law is stated in terms of the color of the cards, but the original sequence supplies only the RANK and SUIT of each card. Table 3 provides some examples of generating laws from Eleusis. Note that the terms in which these laws are expressed (e.g., 'strings of cards of the same suit', 'alternating sequence') are quite different from terms such as RANK and SUIT that described the original sequence. To bridge this difference, appropriate description-space transformations have to be performed.

Eleusis also provides a good domain for studying the use of models for guiding the induction process. The space of possible Eleusis rules using descriptors such as SUIT, RANK, COLOR, FACEDNESS, PARITY, PRIMENESS, and RANK MODULO 3 very large. In our description language, there are more than  $10^{137}$  possible sequence-generating rules involving four or fewer conjunctive expressions.<sup>1</sup> A breadth-first search of this space, such as is conducted by the candidate-elimination algorithm [26], would clearly be impossible. Fortunately, the rules used by people tend to cluster into certain classes that can be well-described by three models: periodic rules, decomposition rules, and disjunctive normal form (DNF) rules. Thus, a model-directed approach can be used to discover sequence-generating rules in Eleusis.

<sup>1</sup>This estimate is based on computing the space of all syntactically legal VL22 conjunctive statements (see Section 5) containing the following set of descriptors (after each descriptor is listed the number of elements in its value set and the number of possible terms that can be formed using those elements): SUIT (4,9), RANK (13,91), COLOR (2,3), FACEDNESS (2,3), PARITY (2,3), PRIMENESS (2,3), RANKMOD3 (3,7), D-SUIT01 (4,9), D-SUIT02 (4,9), D-RANK01 (25,300), D-RANK02 (25,300), S-RANK01 (25,300), S-RANK02 (25,300), D-COLOR01 (2,3), D-COLOR02 (2,3), D-FACEDNESS01 (2,3), D-FACEDNESS02 (2,3), D-PARITY01 (2,3), D-PARITY02 (2,3), D-PRIMENESS01 (2,3), D-PRIMENESS02 (2,3), D-RANKMODS3-01 (3,7), D-RANKMOD3-02 (3,7). The SUIT and RANKMOD3 descriptors are cyclically ordered, while the RANK descriptors are interval descriptors. All others are nominal. In a decomposition rule with a lookback of  $L = 2$ , the first seven descriptors appear three times—once for each card. Hence, the total number of possible conjuncts is  $(9 \times 91 \times 3 \times 3 \times 3 \times 3 \times 7)^3 \times (9 \times 300 \times 300 \times 3 \times 3 \times 3 \times 3 \times 7)^2 = 2.1122 \times 10^{34}$ . If there are four conjuncts in a rule, then we obtain  $[2.1122 \times 10^{34}]^4 = 1.99 \times 10^{137}$ .

TABLE 3. Some examples of sequence-generating rules in Eleusis

---

If the last card was a spade, play a heart; if last card was a heart, play diamonds; if last was diamond, play clubs; and if last was club, play spades.

The card played must be one point higher than or one point lower than the last card.

If the last card was black, play a card higher than or equal to that card; if the last card was red, play lower or equal.

Play alternating even and odd cards.

Play strings of cards such that each string contains cards all in the same suit and has an odd number of cards in it.

---

### 3. Overview of Solution

This section gives an overview of the approach taken to solving the NDP problem defined in Section 2. The approach is a combination of bottom-up data transformation, top-down model specialization, and bidirectional search to fit the specialized models to the transformed data. These three processes can be briefly explained as follows:

(1) Bottom-up data transformation involves applying various transformation operators to the initial sequence description to obtain a *derived sequence description*. We use four basic data transformations: adding derived attributes, segmenting, splitting, and blocking. Details of these are described in Section 4.

(2) Top-down model specialization involves specifying particular values for the parameters of general rule models to obtain a *parameterized model*. We use three general models: the disjunctive normal form model (DNF), the decomposition model, and the periodic model. Each of these models has one or more parameters. For example, both the DNF and decomposition models have a single parameter: the *lookback*,  $L$  (i.e., the number of objects back from the given object in the sequence that are assumed to determine the next object). The periodic model has two parameters: the *lookback*,  $L$ , and the *period length*,  $P$ , which indicates the length of the repeating period in the sequence. Details of the model specialization process are described in Section 5.

(3) The model-instantiation process attempts to fit the parameterized model to the derived sequence description to produce an instantiated parameterized model. A model that has been parameterized and instantiated serves as a sequence-generating rule. This process is described in Section 6.

The above three steps are illustrated schematically in Fig. 3.

Model instantiation, as used in this paper, is an extension of the well-known AI technique of schema instantiation. Schema instantiation has been applied, for example, by Schank and Abelson [31] to interpret natural language, by Englemore and Terry [10] to interpret X-ray diffraction data in protein




FIG. 3. Schematic description of the rule discovery process.

chemistry, and by Friedland [11] to plan genetics experiments. Model instantiation differs from schema instantiation in the complexity of the instantiation process. Model instantiation involves not only filling in predetermined slots or substituting constants for variables, but also synthesizing a logical formula of an assumed type. For example, in order to instantiate each of the three models described below, the program must synthesize a conjunction of predicates or a disjunction of such conjunctions that satisfies certain constraints. Model-instantiation methods share with schema-instantiation methods the advantage that they are efficient and also effective with noisy and uncertain data. The constraints provided by the models (or schemas) drastically reduce the size of the space that the program must search.

The principal disadvantage of model and schema-instantiation methods is that they require substantial amounts of domain knowledge to be built into the program. In order to keep this domain knowledge explicit and easily modified, we employ a ring architecture in the design of the learning program, as described in Section 6. This architecture facilitates the application of the system to a variety of problems by simplifying the process of changing the domain-specific parts of the program.

#### 4. Describing and Transforming Training Instances

Now that we have defined the problem to be solved (the NDP problem) and sketched the solution, we launch into the details of that solution. This section presents the description language for representing the original sequences and the transformation operators that can be applied to modify that representation.

#### 4.1. Representing the initial sequence

A sequence of objects is represented as an indexed sequence<sup>2</sup>

$$\langle q_1, q_2, \dots, q_k \rangle.$$

It is assumed that the only relevant relationship between two objects is their ordering in the sequence. Each object is described by a set of attributes (also called *descriptors*)  $f_1, \dots, f_n$ , which can be viewed as functions mapping objects into attribute values. To state that attribute  $f_i$  of object  $q_j$  has value  $r$ , we write

$$[f_i(q_j) = r].$$

This expression is called a *selector*. For example, if  $f_i$  is *color* and  $r$  is *red*, then the selector

$$[\text{color}(q_j) = \text{red}]$$

states that the color of the  $j$ th object in the sequence is red.

Each attribute is only permitted to take on values from a finite value set called the *domain*,  $D(f_i)$ , of that attribute. This constraint is part of the background knowledge that has to be given to the program. For example, in a deck of cards, the domain of the *suit* attribute is {clubs, diamonds, hearts, spades}. Additional knowledge about the domain set can be represented. In particular, the domain set may be linearly ordered, cyclically ordered (i.e., have a circular, wrap-around ordering), or tree ordered. We will see below how these domain orderings are applied to the problem of representing cards in an Eleusis game.

A complete initial description of a single object,  $q_j$ , called an *event*, is an expression giving the values for all of the attributes of  $q_j$ . This is usually written as a conjunction of selectors:

$$[f_1(q_j) = r_1][f_2(q_j) = r_2] \cdots [f_n(q_j) = r_n].$$

It can also be represented as a vector of attribute values:

$$(r_1, r_2, \dots, r_n).$$

This vector notation suggests that each object description can be viewed as a point in the *event space*  $E$ :

<sup>2</sup>A summary of the notational conventions used in this paper appears in Appendix A.

$$E = D(f_1) \times D(f_2) \times \cdots \times D(f_n)$$

where  $D(f_i)$  is the domain of attribute  $f_i$ . This event space contains all possible events.

A complete description of the initial sequence is a sequence of conjunctions of selectors (or alternatively, a sequence of attribute vectors)—one conjunction for each object in the sequence. Each complete description can be viewed as a trajectory in the event space,  $E$ . The space of all possible sequences can be generated by selecting all possible sequences of events chosen from  $E$ .

## 4.2. Transforming the sequence

As we mentioned in Section 1, it is often necessary to transform the initial sequence into a *derived sequence* in order to facilitate the discovery of sequence-generating rules. Such a data transformation can be viewed as a mapping  $T$  from one set of sequences  $S$ , containing objects  $Q$ , described by attributes  $F$ , to another set of *derived sequences*  $S'$ , containing *derived objects*  $Q'$ , and described by *derived attributes*  $F'$ .

$$T_{p_1, \dots, p_k} : \langle S, Q, F \rangle \rightarrow \langle S', Q', F' \rangle$$

where  $p_1, \dots, p_k$  are parameters of the transformation that control its application. We have found four basic transformations to be especially useful for discovering sequence-generating rules: (a) *adding derived attributes*, (b) *segmenting*, (c) *splitting into phases*, and (d) *blocking*. Each of these is described in turn.

### 4.2.1. Adding derived attributes

The simplest transformation does not change the set of sequences,  $S$ , or the set of objects,  $Q$ , but only the set of attributes,  $F$ . For example, in Eleusis, the initial set  $F$  contains only two attributes: the RANK and SUIT of a card. These can be augmented by deriving such attributes as COLOR (red or black), FACEDNESS (faced or nonfaced), PARITY (odd or even), and PRIMENESS (prime or not prime in rank). Although the adding-derived-attributes transformation has no parameters, in cases where many such attributes could be derived it would be necessary to provide the program with some heuristics to suggest which attributes should be derived and added to the derived sequence.

### 4.2.2. Segmenting

The segmenting transformation derives a new sequence that is made up of a new set of objects,  $Q'$ , which are described by a new set of attributes,  $F'$ . The



new sequence is produced from the original sequence by dividing the original sequence into non-overlapping segments. Each segment becomes a derived object in the new sequence. The only parameter of this transformation is the segmentation condition that specifies how the original sequence is to be divided into segments. Three types of segmentation conditions are distinguished: (a) those that use properties of the *original objects* to determine where the sequence *should* be broken, (b) those that use properties of the *original objects* to determine where the sequence *should not* be broken, and (c) those that use properties of *derived objects* to determine where the original sequence *should* be broken.

For example, suppose the original sequence consists of physical objects described by attributes such as WEIGHT, COLOR, and HEIGHT. An example of each type of segmentation condition follows:

(1) Break when

$$[\text{weight}(q_{i-1}) > 10][\text{weight}(q_i) \leq 10].$$

According to this condition, the original sequence is to be broken (between objects  $q_{i-1}$  and  $q_i$ ) at the point where the weight of an object changes from above 10 to under 10.

(2) Don't break as long as

$$[\text{color}(q_i) = \text{color}(q_{i-1})][\text{weight}(q_i) > 10].$$

This condition states that the original sequence will not be broken (between  $q_{i-1}$  and  $q_i$ ) if the color stays the same and the weight remains above 10. It will be broken at any point where these conditions do not both hold.

(3) Break so that

$$[\text{length}(q_i') = 2].$$

This condition states that derived objects ( $q_i'$ ) should be subsequences of length 2 from the original sequence (i.e., pairs of adjacent objects from the original sequence).

The choice of attributes,  $F'$ , for describing the newly-derived objects,  $Q'$ , depends on the segmentation condition used to segment the sequence. For example, if the segmentation condition is  $[\text{length}(q_i') = 2]$ , attributes of interest might include the sum of the VALUES of the two original objects, the maximum VALUE, the minimum VALUE, and so on. The LENGTH of the segment would not be of interest, since by definition, it is a constant. However, if the segmentation

condition is  $[\text{color}(q_i) = \text{color}(q_{i-1})]$ , the LENGTH of the segment could be an interesting attribute and should be derived. Also, the COLOR shared by all of the cards in the segment might be of interest. In our implementation, the user provides an a priori knowledge base that specifies which attributes should be derived. Every user-specified attribute is derived unless the program can prove from the segmentation condition that the attribute would not have a well-defined value for each segment in the sequence or else would be trivially constant for all segments.

Often, a segmentation condition leads to the creation of incomplete segments at the beginning and end of the original sequence. These boundary cases can create difficulties during model instantiation, so they are ignored during rule discovery, but checked during rule evaluation.

#### 4.2.3. Splitting

The splitting transformation splits a single sequence into a *sequence* of  $P$  separate subsequences called *phases*:  $\langle ph_1, ph_2, \dots \rangle$ . Sequence  $ph_i$  starts with object  $q_i$  (the object at the  $i$ th position in the original sequence) and continues with objects taken from succeeding positions at distance  $P$  apart in the original sequence. The objects in phase sequence  $ph_i$  are referred to as  $\langle ph_{i,1}, ph_{i,2}, ph_{i,3}, \dots \rangle$  or  $\langle phi1, phi2, phi3, \dots \rangle$ . Hence, after splitting, derived object  $ph_{i,j}$  is identical to original object  $q_{i+P \times (j-1)}$ .  $P$  is the parameter of the splitting transformation and denotes the number of phases (the period length). Fig. 4 shows the splitting operation with  $P = 3$ .

Original sequence:	$\langle q_1,$	$q_2,$	$q_3,$	$q_4,$	$q_5,$	$q_6,$	$q_7,$	$q_8,$	$q_9 \rangle$
Derived sequence:	$\langle ph_1,$	$ph_2,$	$ph_3 \rangle,$	where					
$ph_1:$	$\langle ph_{11}$			$ph_{12}$			$ph_{13} \rangle$		
$ph_2:$		$\langle ph_{21}$			$ph_{22}$			$ph_{23} \rangle$	
$ph_3:$			$\langle ph_{31}$			$ph_{32}$			$ph_{33} \rangle$

FIG. 4. Splitting transformation with  $P = 3$ .

The objects within each phase retain the linear ordering that they had in the original sequence. The phases themselves can be considered to be cyclically ordered so that  $ph_1$  precedes  $ph_2$ , which precedes  $ph_3$ , and so on, until  $ph_P$ , which is followed by  $ph_1$  again. Consider, for example, the following sequence:

$$\langle 1, 8, 2, 9, 3, 10, 4, 11 \rangle.$$

The splitting transformation with  $P = 2$  would produce the sequence  $\langle ph_1, ph_2 \rangle$  where

$$ph_1 = \langle 1, 2, 3, 4 \rangle, \quad ph_2 = \langle 8, 9, 10, 11 \rangle.$$

Since the splitting transformation simply breaks the original sequence of

objects into subsequences of the same objects, no new descriptors are defined. The descriptors used to characterize objects in each of the phases are the same as those used to characterize the objects in the original sequence.

The splitting transformation can be applied to break one sequence-prediction problem into several subproblems—one for each phase. This enables the system to discover periodic rules.

#### 4.2.4. Blocking

The blocking transformation converts the original sequence into a new sequence made up of a new set of objects  $B'$  and a new set of attributes  $F'$ . The new sequence is created by breaking the original sequence into overlapping segments called *blocks*. Each object  $b_i$  in the new sequence describes a block of  $L + 1$  consecutive objects from the original sequence, starting at object  $q_i$  (called the *head*) and proceeding backwards to object  $q_{i-L}$  (where  $L$  is the *lookback* parameter of the blocking transformation). Fig. 5 shows the blocking operation for  $L = 2$  (block length of 3).

Several attributes are derived to describe each block. For each attribute  $A$  applicable to the objects in the original sequence, the attributes  $A_0, A_1, \dots, A_L$  are defined that are applicable to the objects in the derived sequence.  $A_0(b_i)$  has the same value as  $A(q_i)$ ;  $A_1(b_i)$  has the same value as  $A(q_{i-1})$ ; and so on until  $A_L(b_i)$ , which has the same value as  $A(q_{i-L})$ . In other words, the original attributes are retained in the new sequence, but they are renamed so that they apply to whole *blocks* rather than to individual objects in the original sequence. The numerical suffix on the new names encodes the relative position of the original object  $q_i$  in block  $b_j$ .

Original sequence:	⟨q1,	q2,	q3,	q4,	q5,	q6,	q7,	q8⟩
Derived sequence:			⟨b3,	b4,	b5,	b6,	b7,	b8⟩
where $b_i$ are derived objects defined as follows:								
b3:	⟨q1,	q2,	<u>q3</u> ⟩					
b4:		⟨q2,	<u>q3</u> ,	<u>q4</u> ⟩				
b5:			⟨q3,	<u>q4</u> ,	<u>q5</u> ⟩			
b6:				⟨q4,	<u>q5</u> ,	<u>q6</u> ⟩		
b7:					⟨q5,	<u>q6</u> ,	<u>q7</u> ⟩	
b8:						⟨q6,	<u>q7</u> ,	<u>q8</u> ⟩

The underlined objects are the head objects of each block.

FIG. 5. The blocking transformation with lookback parameter  $L = 2$ .

For example, suppose the original sequence of objects is  $\langle q1, q2, q3, q4, q5 \rangle$  with attributes *RANK* and *SUIT*, where

$$\begin{aligned}
 &[\text{rank}(q1) = 2][\text{suit}(q1) = H], \\
 &[\text{rank}(q2) = 4][\text{suit}(q2) = S], \\
 &[\text{rank}(q3) = 6][\text{suit}(q3) = C], \\
 &[\text{rank}(q4) = 8][\text{suit}(q4) = D], \\
 &[\text{rank}(q5) = 10][\text{suit}(q5) = H].
 \end{aligned}$$

Suppose we apply the blocking transformation to this sequence with  $L = 2$  to obtain the derived sequence of blocks  $\langle b_3, b_4, b_5 \rangle$ . Then the descriptors RANK0, RANK1, RANK2, SUIT0, SUIT1, and SUIT2 will be derived with the values

$$\begin{aligned} &[\text{rank2}(b_3) = 2][\text{suit2}(b_3) = H] \ \& \\ &[\text{rank1}(b_3) = 4][\text{suit1}(b_3) = S] \ \& \\ &[\text{rank0}(b_3) = 6][\text{suit0}(b_3) = C] \end{aligned}$$

$$\begin{aligned} &[\text{rank2}(b_4) = 4][\text{suit2}(b_4) = S] \ \& \\ &[\text{rank1}(b_4) = 6][\text{suit1}(b_4) = C] \ \& \\ &[\text{rank0}(b_4) = 8][\text{suit0}(b_4) = D] \end{aligned}$$

$$\begin{aligned} &[\text{rank2}(b_5) = 6][\text{suit2}(b_5) = C] \ \& \\ &[\text{rank1}(b_5) = 8][\text{suit1}(b_5) = D] \ \& \\ &[\text{rank0}(b_5) = 10][\text{suit0}(b_5) = H]. \end{aligned}$$

This transformation leads to a highly redundant representation of the information in the original sequence. For example, the information about SUIT and RANK of the original object  $q_3$  is repeated as SUIT0 and RANK0 of block  $b_3$ , SUIT1 and RANK1 of block  $b_4$ , and SUIT2 and RANK2 of block  $b_5$ . However, this derived sequence of blocks facilitates the representation of the relationships between objects in the original sequence. Many sequence-prediction rules involve such relationships.

To represent relationships between objects, additional descriptors called *sum* and *difference* descriptors are defined. In the case of the above sequence, the descriptors S-RANK01, S-RANK02, D-RANK01, D-RANK02, D-SUIT01, and D-SUIT02 are created. The value of S-RANK01( $b_i$ ) is the sum of RANK0( $b_i$ ) and RANK1( $b_i$ ). The value of D-RANK01( $b_i$ ) is the difference between RANK0( $b_i$ ) and RANK1( $b_i$ ). Thus, in addition to the selectors shown above, the following selectors would also be derived for the new sequence:

$$\begin{aligned} &[\text{s-rank01}(b_3) = 10][\text{s-rank02}(b_3) = 8] \ \& \\ &[\text{d-rank01}(b_3) = 2][\text{d-rank02}(b_3) = 4] \ \& \\ &[\text{d-suit01}(b_3) = 1][\text{d-suit02}(b_3) = 2] \end{aligned}$$

$$\begin{aligned} &[\text{s-rank01}(b_4) = 14][\text{s-rank02}(b_4) = 12] \ \& \\ &[\text{d-rank01}(b_4) = 2][\text{d-rank02}(b_4) = 4] \ \& \\ &[\text{d-suit01}(b_4) = 1][\text{d-suit02}(b_4) = 2] \end{aligned}$$

$$\begin{aligned} &[\text{s-rank01}(b_5) = 18][\text{s-rank02}(b_5) = 16] \ \& \\ &[\text{d-rank01}(b_5) = 2][\text{d-rank02}(b_5) = 4] \ \& \\ &[\text{d-suit01}(b_5) = 1][\text{d-suit02}(b_5) = 2]. \end{aligned}$$

Using this representation, it is relatively easy to discover that  $[\text{d-rank01}(b_i) = 2]$  is true for all blocks  $b_i$ .

Ordinarily, sum and difference attributes only make sense for attributes such as RANK whose domain sets are linearly ordered. We have extended the definition of difference to cover unordered and cyclically ordered domain sets as well. For an unordered attribute such as COLOR, whose domain set is {red, black}, D-COLOR<sub>01</sub> takes on the value 0 if the COLOR<sub>0</sub>( $b_i$ ) = COLOR<sub>1</sub>( $b_i$ ) and 1 otherwise. For attributes with cyclically ordered domain sets, such as SUIT (with values {clubs, diamonds, hearts, spades}), D-SUIT<sub>01</sub> is equal to the number of steps in the forward direction that are required to get from SUIT<sub>1</sub>( $b_i$ ) to SUIT<sub>0</sub>( $b_i$ ). If SUIT<sub>1</sub>( $b_i$ ) = diamonds and SUIT<sub>0</sub>( $b_i$ ) = clubs, D-SUIT<sub>01</sub>( $b_i$ ) = 3.

The sum and difference attributes make the ordering of the original sequence explicit in the attributes that describe each block. Consequently, it is no longer necessary to represent the ordering between blocks. Hence, the model-fitting algorithms discussed below treat the derived sequence (of blocks) as an unordered set of events.

One difficulty with the above notation is that the numerical suffixes are not very easy to read, especially when they are combined with sum or difference prefixes. Hence, we have developed an alternative representation that is more comprehensible. In this notation, selectors that refer to blocks, such as [suit1( $b_i$ ) = H], are written as selectors that refer to objects in the original sequence, such as [suit( $q_i - 1$ ) = H]. Similarly, selectors such as [d-rank01( $b_i$ ) = 3] are written as [rank( $q_i$ ) = rank( $q_i - 1$ ) + 3]. This notation makes the meaning of the selectors clear without having to explicitly mention the blocks  $b_i$ . For purposes of implementation, the first notation is better because it enables the program to treat all sequences—including derived sequences—uniformly. In contrast to this, the second notation only works when at most one blocking transformation is applied. Multiple blocking transformations cannot be captured without explicitly mentioning the blocks. Since it is rare that more than one blocking transformation is needed and since the second notation is more understandable, we will use it for the rest of this paper.

### 5. Representing Sequence-generating Rules and Models

A *sequence-generating rule* is a function  $g$  that assigns to each sequence of objects,  $\langle q_1, q_2, \dots, q_k \rangle$ , a non-empty set of *admissible next objects*  $Q_{k+1}$ :

$$g : \{ \langle q_1, q_2, \dots, q_k \rangle \} \rightarrow \{ Q_{k+1} \},$$

$Q_{k+1}$  is the set of all objects that could appear as the next object in the sequence. For example, in the rule 'Play a card whose rank is one higher than the previous card', the value of the function  $g$  when applied to a sequence whose last card is the 4C is the set of cards {5C, 5D, 5H, 5S}. This is written

$$g(\langle \dots, 4C \rangle) = Q_{k+1} = \{5C, 5D, 5H, 5S\}.$$

Each set  $Q_{k+1}$  may contain only one event, or it may contain a large set of

possible events. If for all  $k$ , the sequence  $\langle q_1, q_2, \dots, q_k \rangle$  is mapped by  $g$  into a singleton set, then the rule is a *deterministic* rule; otherwise, it is a *nondeterministic* rule. This paper addresses the problem of discovering a nondeterministic sequence-generating rule,  $g$ , given the sequence  $\langle q_1, q_2, \dots, q_k \rangle$ , where the  $q_i$  are characterized by a finite set of discrete attributes.

The sequence  $\langle q_1, q_2, \dots, q_k \rangle$  can be viewed as the set of assertions

$$\begin{aligned} q_1 &\in g(\langle \rangle), \\ q_2 &\in g(\langle q_1 \rangle), \\ &\vdots \\ q_m &\in g(\langle q_1, \dots, q_{m-1} \rangle). \end{aligned}$$

(Recall that the value of  $g(s)$ , where  $s$  is a sequence, is a set of possible next objects.) These assertions are positive instances of the desired sequence-generating rule.

In Eleusis, negative instances are provided by the cards on the side lines—that is, the cards indicated by the dealer as being incorrect continuations of the sequence. A side-line card  $q_4^-$  played after card  $q_3$  provides a negative instance of the form:

$$q_4^- \notin g(\langle q_1, q_2, q_3 \rangle).$$

The goal is to find a description for  $g$  that is *consistent* with these training instances and satisfies some preference criterion. (A description  $g$  is consistent with the training instances if it characterizes all of the positive instances and excludes all of the negative instances.)

The preference criterion in our methodology, and generally in learning systems, attempts to evaluate a candidate rule in terms of its generality, predictive power, simplicity, and so on. These semantic properties are difficult to compute, however. Instead, virtually all learning systems employ syntactic criteria that correspond in some way to these semantic criteria. Syntactic criteria—such as the number of selectors in a conjunction and the number of conjuncts in a disjunction—will only correspond to the semantic criteria if the representational framework is well chosen (See [18]). As we noted in the introduction, most previous AI research on learning has employed a single representational framework or model for describing the rules or concepts to be learned. In Eleusis, a single framework is insufficient. Instead, we have developed three basic models that were found to be useful: the DNF model, the decomposition model, and the periodic model. When these models are employed, syntactic criteria can be used to approximate semantic criteria during evaluation.

A *model* is a structure that specifies the syntactic form of a class of descriptions (in our case, sequence-generating rules). A model consists of

*model parameters* and a set of *constraints* that the model places on the forms of descriptions. The process of specifying the values for the parameters of a model is called *parameterizing* the model. The process of filling in the form of the parameterized model is called *instantiating* the model. A fully-parameterized and fully-instantiated model forms a sequence-generating rule. Models can be instantiated using the original sequence, or, more typically, using a sequence derived by applying some of the data transformations discussed in the previous section.

All three models use the variable-valued logic calculus VL22 for representing sequence-generating rules. VL22 is an extension of the predicate calculus that uses the *selector* as its simplest kind of formula. The VL22 selector is substantially more expressive than the simple selector presented above in Section 4.1. Recall that the simple selector has the form

$$[f_i(q_j) = r]$$

whereas the VL22 selector has the form

$$[f_i(x_1, x_2, \dots, x_n) = r_1 \vee r_2 \vee \dots \vee r_m].$$

In the VL22 selector, attributes  $f_i$  can take any number of arguments  $(x_1, x_2, \dots, x_n)$ . Furthermore, the attributes  $f_i$  can take on any one of a *set* of values  $\{r_1, r_2, \dots, r_m\}$ . The  $\vee$  denotes the *internal disjunction* operator. Thus, the selector

$$[\text{rank}(q_i) = 9 \vee 10 \vee J \vee Q \vee K]$$

indicates that the rank of object  $q_i$  can be either 9, 10, J, Q, or K. The internal disjunction represents disjunction over the values of a single variable. In this case, it could be expressed alternatively as

$$[\text{rank}(q_i) \geq 9],$$

since the domain of the RANK attribute is known to be linearly ordered with a maximum value of K (King). To aid comprehensibility, VL22 provides the operators  $<$ ,  $>$ ,  $\leq$ ,  $\geq$ , and  $\neq$ , in addition to the basic  $=$  operator.

Examples of typical selectors include:

- $[\text{rank}(q_i) \neq \text{rank}(q_{i-1})]$  (paraphrase: the RANK of  $q_i$  is different from the RANK of  $q_{i-1}$ );
- $[\text{suit}(q_i) = \text{suit}(q_{i-1}) + 1]$  (paraphrase: the SUIT increases by one from  $q_{i-1}$  to  $q_i$ );
- $[\text{rank}(q_i) + \text{rank}(q_{i-2}) > 10]$  (paraphrase: the sum of the RANKs of  $q_i$  and  $q_{i-2}$  is greater than 10).

Now that we have introduced the basic notation of VL22, each of the three rule models is presented in turn.

### 5.1. The DNF model

The DNF model supports the broad class of rules that can be expressed as a universally quantified VL22 statement in disjunctive normal form. The DNF model has one parameter, the degree of lookback,  $L$ . An example of a DNF rule (with  $L = 1$ ) is

$$\forall i ([\text{color}(q_i) = \text{color}(q_{i-1})] \vee [\text{rank}(q_i) = \text{rank}(q_{i-1})])$$

which can be paraphrased as "Every object ( $q_i$ ) in the sequence has the same color or the same rank (or both) as the preceding object ( $q_{i-1}$ )."

In general, a DNF rule is a collection of conjuncts of the form

$$\forall i (C_1 \vee C_2 \vee C_3 \vee \dots \vee C_k).$$

The universal quantification over  $i$  indicates that this description is true for all objects  $q_i$  in the sequence.

An additional constraint specified in the DNF model is that the number of conjuncts,  $k$ , should be close to the minimum that produces a description consistent with the data.

### 5.2. The decomposition model

The decomposition model constrains the description to be a set of implications of the form:

$$L_1 \Rightarrow R_1, \quad L_2 \Rightarrow R_2, \quad \dots, \quad L_m \Rightarrow R_m,$$

where the  $\Rightarrow$  sign indicates logical implication.

The model states that the left- and right-hand sides,  $L_j$  and  $R_j$ , must all be VL22 conjunctions. The left-hand sides must be mutually exclusive and exhaustive—that is,

$$L_1 \vee L_2 \vee \dots \vee L_m \quad \text{and} \quad \forall j, k (j \neq k \Rightarrow \neg (L_j \wedge L_k)).$$

A decomposition rule describes the next object in the sequence in terms of characteristics of the previous objects in the sequence. For example, the rule

$$\forall i ([\text{color}(q_{i-1}) = \text{black}] \Rightarrow [\text{parity}(q_i) = \text{odd}] \ \& \\ [\text{color}(q_{i-1}) = \text{red}] \Rightarrow [\text{parity}(q_i) = \text{even}])$$

is a decomposition rule that says that if the last card was black, the next card



must be odd, and if the last card was red, the next card must be even.

The decomposition model has a lookback parameter,  $L$ , that indicates how far back in the sequence the sequence-generating rule must 'look' in order to predict the next object in the sequence. The above rule has a lookback parameter of 1, because it examines  $q_{i-1}$  (the previous object in the sequence).

### 5.3. The periodic model

This model consists of rules that describe objects in the sequence as having attribute values that repeat periodically. For example, the rule "Play alternating red and black cards" is a periodic rule. The periodic model has two parameters: the period length,  $P$ , and the lookback,  $L$ . The period length parameter,  $P$ , gives the number of phases in the periodic rule. A periodic rule can be viewed as applying a splitting transformation to split the original sequence into  $P$  separate sequences. Each separate phase sequence has a simple description. The lookback parameter,  $L$ , tells how far back, within a phase sequence, a periodic rule 'looks' in order to predict the attributes of the next object in that phase. The periodic model imposes the additional constraint (or preference) that the different phases be disjoint (i.e., any given card is only playable within one phase).

A periodic rule is represented as an ordered  $P$ -tuple of VL22 conjunctions. The  $j$ th conjunct describes the  $j$ th phase sequence. The rule

$$\langle [\text{color}(\text{ph}1i) = \text{red}], [\text{rank}(\text{ph}2i) \geq \text{rank}(\text{ph}2i - 1)] \rangle$$

is a periodic rule with  $P = 2$  and  $L = 1$ , which says that the sequence is made of two (interleaved) phases. Each card in the first phase is red; each card in the second phase has at least as large a rank as the preceding card in that phase. Hence, one sequence that satisfies this rule is  $\langle 2H, 3C, 10H, 5S, AD, 6S, 6H, 6C \rangle$ .

A more complex periodic rule is the one used to generate the sequence shown in Fig. 1. It can be represented as

$$\langle [\text{texture-of-nodes}(\text{ph}1i) = \text{solid black}] \& \\ [\text{shape}(\text{ph}1i) = \text{T-junction}] \& \\ [\text{orientation}(\text{ph}1i) = \text{orientation}(\text{ph}1i - 1) - 45],$$

$$[\text{texture-of-nodes}(\text{ph}2i) = \text{clear}] \& \\ \langle [\text{number-of-nodes}(\text{ph}2\text{ph}1i) = 4], \\ [\text{number-of-nodes}(\text{ph}2\text{ph}2i) = 8] \rangle,$$

$$[\text{texture-of-nodes}(\text{ph}3i) = \text{cross}] \& \\ [\text{shape}(\text{ph}3i) = \text{bar}] \& \\ [\text{orientation}(\text{ph}3i) = \text{orientation}(\text{ph}3i - 1) + 45] \rangle$$

Notice that this is a periodic rule with three phases and a lookback of 1. The middle phase of the period is itself a periodic rule with the NUMBER-OF-NODES alternating between 4 and 8.

#### 5.4. Derived models

The three basic models can be combined to describe more complex rules. Basic models can be joined by conjunction, disjunction, and negation. For example, the rule "play alternating red and black cards such that the cards are in nondecreasing order" is a conjunction of the periodic rule

$$\langle [\text{color}(\text{ph}1i) = \text{red}], [\text{color}(\text{ph}2i) = \text{black}] \rangle$$

and the DNF rule

$$[\text{rank}(q_i) \geq \text{rank}(q_i - 1)].$$

#### 5.5. Model equivalences and the heuristic value of models

The reader may have noticed that the decomposition and periodic models appear to be special cases of the DNF model. In particular, assuming that the left-hand side clauses in a decomposition rule are mutually-exclusive and exhaustive, the decomposition rule

$$L_1 \Rightarrow R_1 \&$$

$$L_2 \Rightarrow R_2 \&$$

$$\vdots$$

$$L_m \Rightarrow R_m$$

can be written as the DNF rule

$$[L_1 \& R_1] \vee [L_2 \& R_2] \vee \cdots \vee [L_m \& R_m].$$

Similarly, if the phases ( $C_i$ ) of a periodic rule are mutually-exclusive and exhaustive, then the periodic rule

$$\langle C_1, C_2, \dots, C_k \rangle$$

can be reexpressed as a decomposition rule of the form

$$C_1 \Rightarrow C_2 \&$$

$$C_2 \Rightarrow C_3 \&$$

$$\vdots$$

$$C_{k-1} \Rightarrow C_k \&$$

$$C_k \Rightarrow C_1.$$

This transformation from the periodic model into the decomposition model

does not work when the phases of the periodic rule overlap. Consider, for example, the following periodic rule in which the phases are neither mutually exclusive nor exhaustive:

$$\langle (\text{color}(\text{ph}1i) = \text{red}), [\text{rank}(\text{ph}2i) = \text{even}] \rangle$$

(paraphrase: play alternating red and even cards).

Because the phases overlap, the transformation into a decomposition rule produces a slightly different rule:

$$\begin{aligned} &[\text{color}(q_i - 1) = \text{red}] \Rightarrow [\text{parity}(q_i) = \text{even}] \ \& \\ &[\text{parity}(q_i - 1) = \text{even}] \Rightarrow [\text{color}(q_i) = \text{red}] . \end{aligned}$$

To see how the two rules differ, consider the sequence of cards

$$\langle 3D, 2D, 4C, \dots \rangle .$$

This sequence satisfies the second rule (the first if-then clause can be applied twice), but not the first rule (since the 4C is not red).

Even when the constraints of mutual exclusion and exhaustion are violated, it is always possible to develop some equivalent DNF rule for any periodic or decomposition rule. This is so because one can always provide additional descriptors that capture the particular relationship. The resulting DNF rules are not always as succinct or comprehensible, however, as the same rule expressed using one of the other models.

Consider this same periodic rule. Suppose that a new descriptor, called POSITION, is defined whose value for each object  $q_i$  is the position  $i$  of the object in the sequence. With this descriptor, the above rule can be encoded as

$$\begin{aligned} &[\text{position}(q_i) = \text{odd}] \Rightarrow [\text{color}(q_i) = \text{red}] \ \& \\ &[\text{position}(q_i) = \text{even}] \Rightarrow [\text{parity}(q_i) = \text{even}] . \end{aligned}$$

Since any sequence-generating rule can be expressed in the DNF model, it is reasonable to ask why multiple models should be used at all. The answer is that the primary value of multiple models is that they provide heuristic guidance to the search for *plausible* rules. Hence, though the DNF model is capable of *representing* all of these rules, it is not helpful for *discovering* them. In short, it is *epistemologically adequate* but not *heuristically adequate* (see [19, 20]). Each model directs the attention of the learning system to a small subspace of the space of all possible DNF VL22 rules. The next section shows how the constraints associated with each model are incorporated into special model-fitting induction algorithms.

## 6. Architecture and Algorithms

Section 3 described the three basic processes involved in discovering sequence-generating rules: (a) transformation of the original sequence into a derived sequence, (b) selection of an appropriate model (and determination of its parameters) for the given sequence, and (c) fitting (instantiation) of the models to the derived sequence. Sections 4 and 5 presented the four data transformations and the three models. This section covers the third step of fitting specialized models to the transformed sequence. The model-fitting process is most easily understood in the context of the program architecture, so this section also discusses the system's architecture.

### 6.1. Overview of the system

The processes in the system (see Fig. 6) are structured into four components—the three basic ones mentioned above plus an evaluation component. The processes of transforming the initial sequence and of selecting and parameterizing a model are performed in parallel. Then, model-fitting algorithms use the transformed sequence to instantiate the parameterized model to obtain a candidate sequence-generating rule. These candidate rules are then evaluated to determine the final set of rules.

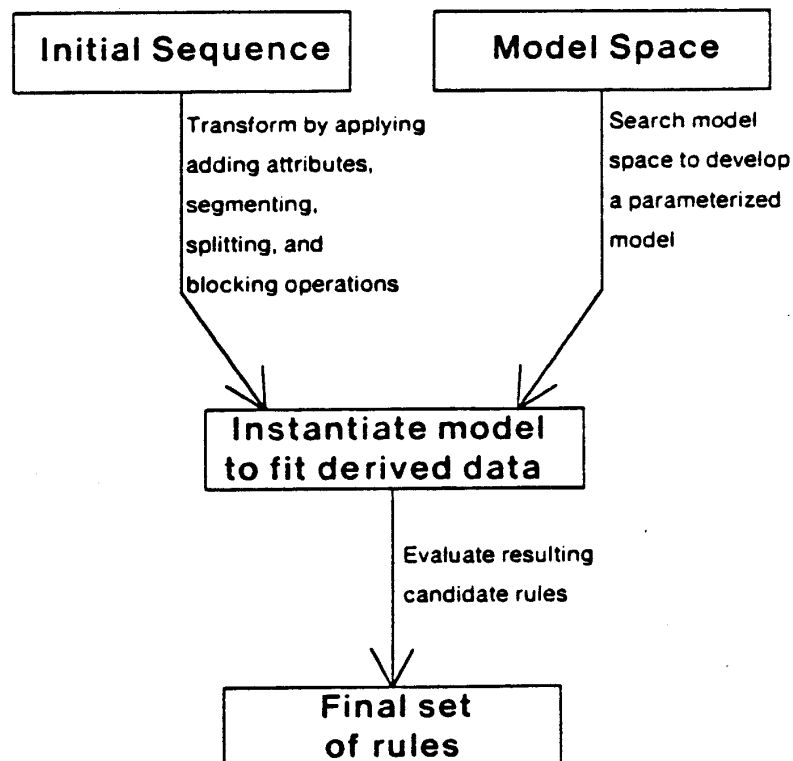


FIG. 6. The model-fitting approach.

The reason for performing data transformation and model selection in parallel is that these two processes are interdependent. For example, if a periodic model is selected (with period length  $P$ ), then a splitting transformation (with number of phases  $P$ ) needs to be applied to the sequence. These two processes can be viewed as simultaneous cooperative searches of two spaces: the space of possible data transformations and the space of possible parameterized models.

## 6.2. Overview of the concentric ring architecture

In order for the learning system to be easily modified to handle entire classes of NDP problems, the system is structured as a set of concentric knowledge rings (see Fig. 7). A knowledge ring is a set of routines that perform a particular function using only knowledge appropriate to that function. The procedures within a given ring may invoke other procedures in that ring or in rings that are inside the given ring. Under these constraints, the concentric ring structure forms a hierarchically organized system.

Ideally, the rings should be organized so that the outermost ring uses the most problem-specific knowledge and performs the most problem-specific operations and the inner-most ring uses the most general knowledge and performs the most general tasks. Such an architecture improves the program's generality because it can be applied to increasingly different NDP problems by removing and replacing the outer rings. In order to apply the program to radically different learning problems, all but the inner-most ring may need to be replaced.

The ring architecture is used here as follows. The outer-most rings perform user-interface functions and convert the initial sequence from whatever domain-specific notation is being used into a sequence of VL22 events. The inner-most ring performs the model-fitting functions. It expects the data to be properly transformed so that the data have the same form as the models to which they are to be fitted. The intervening rings conduct the simultaneous processes of developing a properly parameterized model and transforming the input sequence into an appropriate derived sequence.

The intervening rings also evaluate the rules discovered by the inner-most ring using the knowledge available in each ring.

## 6.3. The system SPARC (Eleusis version)

SPARC (Sequential PAttern ReCognition) stands for a general program designed to solve a variety of nondeterministic prediction problems using the ring architecture. So far, we have implemented only a more specific version of the program, called SPARC/E, tailored specifically to the problem of rule discovery in the game Eleusis. SPARC/E is made up of five rings, as shown in Fig. 7.

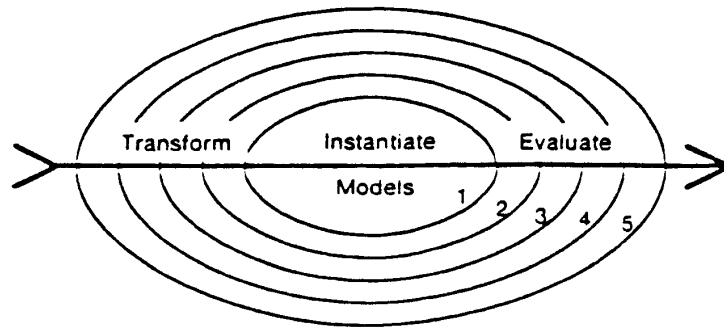


FIG. 7. The knowledge ring architecture.

This section describes the functions of each ring in SPARC/E. To illustrate these ring functions, we use the Eleusis layout shown in Fig. 8. Recall that in an Eleusis layout, the main line shows the correctly-played sequence of cards (positive examples). The side lines, which branch out below the main line, contain cards that do not satisfy the rule—that is, incorrect continuations of the sequence (negative examples).

### 6.3.1. Ring 5: user interface

Ring 5, the outer-most ring, provides a user interface to the program. It executes user's commands for playing the card game Eleusis, as well as commands for controlling the search, data transformation, generalization, and evaluation functions of the program. One command in Ring 5 is the `INDUCE` command that instructs SPARC/E to look for plausible NDP rules that describe the current sequence. When the `INDUCE` command is given, Ring 5 calls Ring 4 to begin the rule discovery process. Ring 5 provides Ring 4 with an initial sequence of VL22 events in which the only attributes are `SUIT` and `RANK`.

Main line:	3H	9S	4C	JD	2C	10D	8H	7H	2C
Side lines:		JD		AH	AS			10H	
		5D		8H	10S				
				QD					

FIG. 8. Sample Eleusis layout.

### 6.3.2. Ring 4: adding derived attributes

Ring 4 applies the adding-derived-attributes transformation to the initial sequence of cards. This involves creating derived attributes that make explicit certain commonly known characteristics of playing cards that are likely to be used in an Eleusis rule: `COLOR`, `PARITY`, `FACED` versus `NON-FACED` cards, and so on. The user of the program provides a definition for each descriptor that is to be derived. Fig. 9 shows the layout from Fig. 8 after it has been processed by Ring 4. The plusses and minuses along the right-hand side of the figure indicate whether

VL22 event	Positive or negative
[rank(card1) = 3][suit(card1) = H] & [parity(card1) = odd][color(card1) = red] & [prime(card1) = N][faced(card1) = Y]	+
[rank(card2) = 9][suit(card2) = S] & [parity(card2) = odd][color(card2) = black] & [prime(card2) = N][faced(card2) = N]	+
[rank(card3) = J][suit(card3) = D] & [parity(card3) = odd][color(card3) = red] & [prime(card3) = Y][faced(card3) = Y]	-
[rank(card3) = 5][suit(card3) = D] & [parity(card3) = odd][color(card3) = red] & [prime(card3) = N][faced(card3) = Y]	-
[rank(card3) = 4][suit(card3) = C] & [parity(card3) = even][color(card3) = black] & [prime(card3) = N][faced(card3) = N]	+
[rank(card4) = J][suit(card4) = D] & [parity(card4) = odd][color(card4) = red] & [prime(card4) = Y][faced(card4) = Y]	+

FIG. 9. Derived layout after Ring 4 processing of the layout in Fig. 8.

the event is a positive example or a negative example of the sequence-generating rule. These derived events are passed to Ring 3 for further processing.

### 6.3.3. Ring 3: segmenting the layout

Ring 3 is the first Eleusis-independent ring. It applies the segmenting transformation to the sequence supplied by Ring 4. In the present implementation, the end points of each segment are determined by applying a segmentation predicate,  $P(\text{card}_{i-1}, \text{card}_i)$  to all pairs of adjacent events in the sequence. When the predicate  $P$  evaluates to FALSE, the sequence is broken between  $\text{card}_{i-1}$  and  $\text{card}_i$  to form the end of a segment. Typical segmentation predicates used are:

[rank(card<sub>i</sub>) = rank(card<sub>i-1</sub>)],  
 [rank(card<sub>i</sub>) = rank(card<sub>i-1</sub>) + 1],  
 [color(card<sub>i</sub>) = color(card<sub>i-1</sub>)],  
 [suit(card<sub>i</sub>) = suit(card<sub>i-1</sub>)],  
 [parity(card<sub>i</sub>) = parity(card<sub>i-1</sub>)].

The layout:

```

3H 5D 2D 7C AC 9C JH 6H 8H QH KS
      5S 4D      AH
                7S

```

The derived sequence:

Description of derived object	Positive or negative
[suit(segment1) = H][color(segment1) = red] & [length(segment1) = 1]	+
[suit(segment2) = D][color(segment2) = red] & [length(segment2) = 2]	+
[suit(segment2) = D][color(segment2) = red] & [length(segment2) = 3]	-
[suit(segment3) = C][color(segment3) = black] & [length(segment3) = 3]	+
[suit(segment4) = H][color(segment4) = red] & [length(segment4) = 4]	+

FIG. 10. Sample layout and segmented sequence under segmentation condition [suit(cardi) = suit(cardi - 1)].

Other techniques for performing segmentation, such as providing a predicate that becomes TRUE at a segment boundary (see Section 4.2.2), are not implemented in SPARC/E.

Ring 3 searches the space of possible segmentations using two search pruning heuristics. After each attempt to segment the sequence, it counts the number of derived objects (segments),  $k$ , in the derived sequence. If  $k$  is less than 3, the segmentation is discarded since there are too few derived objects to use for generalization. If  $k$  is more than half of the number of objects in the original sequence, the segmentation is also discarded because in this case many segments contain only one original object. Segmented sequences that survive these two pruning heuristics are passed on to Ring 2 for further processing.

One segmentation that Ring 3 always performs is the 'null' segmentation—that is, it always passes the unsegmented sequence directly to the inner rings. Fig. 10 shows a sample layout and the resulting derived layout after segmentation using the segmentation condition:

$$[\text{suit}(\text{cardi}) = \text{suit}(\text{cardi} - 1)].$$

The derived objects (segments) are denoted by variables segment $i$ . The negative event



$$[\text{suit}(\text{segment}2) = \text{D}][\text{color}(\text{segment}2) = \text{red}][\text{length}(\text{segment}2) = 3]$$

is obtained from the segment  $\langle 5D, 2D, 4D \rangle$ , which ends in a 'side-line' card. Notice that the very last card in the sequence, the king of spades, is not included in any segment. This is because the king is the first card of a new segment, and it is impossible to know how long that segment will be until it is completed. Once a sequence-generating rule is found by the inner rings, Ring 3 will check to make sure that the king of spades is consistent with the rule.

SPARC/E derives the descriptors COLOR, SUIT, and LENGTH to describe each derived object. The choice of which descriptors to derive involves three steps. First, LENGTH is derived whenever the segmentation transformation is applied. Second, any descriptor that is tested in the segmentation predicate (in this case, SUIT) is also derived. Third, any descriptor is derived whose value can be proved to be the same for all cards in each segment. In this case, COLOR is derived because, if SUIT is a constant, then COLOR is also a constant. Using this segmentation, SPARC can use the DNF model to discover that the segmented sequence can be described as

$$[\text{length}(\text{segment}i) = \text{length}(\text{segment}i - 1) + 1].$$

That is, the LENGTH of each segment of constant SUIT (in the main line) increases by 1.

#### 6.3.4. Ring 2: parameterizing the models

Ring 2 searches the space of parameterizations of the three basic models. Each model is considered in turn. For each model, Ring 2 develops a set of derived events based on each allowed value of the lookback parameter,  $L$ , and the number of phases parameter,  $P$ . The user can control which models should be inspected and what range of values for  $L$  and  $P$  should be investigated. By default, the program will inspect the decomposition model with  $L = 0, 1, \text{ or } 2$ , and the periodic model with  $P = 1 \text{ or } 2$  and  $L = 0 \text{ or } 1$ . The DNF model is not inspected under the default settings for the program.

Specifically, Ring 2 performs the following actions depending on which model is being parameterized:

(A) For the *decomposition* model with lookback parameter  $L$ , Ring 2 applies the blocking transformation to break the sequence received from Ring 3 into blocks of length  $L$ . After blocking, all of the attributes that described the original objects are converted into attributes that describe the whole block (as discussed in Section 4 above). Furthermore, sum and difference descriptors are derived to represent the relationships between adjacent objects in the original sequence. The resulting derived events can be viewed as very specific if-then clauses of the following form.

Given an initial sequence of objects  $\langle q_1, q_2, \dots, q_m \rangle$ , let us look at block  $b_i$

which describes the subsequence  $\langle q_{i-L}, \dots, q_{i-1}, q_i \rangle$ . Let  $F_j$ ,  $j = 0, 1, \dots, L$ , denote the set of selectors describing object  $q_{i-j}$  renamed so that they apply to block  $b_i$ . For example,  $F_1$  could be the selectors

$$[\text{suit1}(\text{bi}) = \text{H}][\text{rank1}(\text{bi}) = 3]$$

—selectors that originally referred to object  $q_{i-1}$ . Let  $d(F_j, F_k)$  denote all of the difference selectors obtained by 'subtracting' event  $F_k$  from event  $F_j$ , and let  $s(F_j, F_k)$  denote all of the summation selectors obtained by 'summing' events  $F_j$  and  $F_k$ . For example,  $d(F_0, F_1)$  could include the selectors

$$[\text{d-suit01}(\text{bi}) = 2][\text{d-rank01}(\text{bi}) = -3]$$

obtained from 'subtracting'  $F_1$  from  $F_0$ .

With these definitions, the derived events for the decomposition model have the form:

$$F_1 \& \dots \& F_L \Rightarrow F_0 \& d(F_0, F_1) \& \dots \& d(F_0, F_L) \& s(F_0, F_1) \& \dots \& s(F_0, F_L).$$

Suppose, for example, that the initial sequence of cards is

$$\langle 2\text{H}, 4\text{D}, 6\text{S}, 8\text{C} \rangle.$$

with only the SUIT and RANK descriptors being employed. Then suppose that Ring 2 applies the blocking transformation with a lookback of 2. Fig. 11 shows the two derived events that will be produced by Ring 2 (the corresponding notation is shown to the right of each group of selectors).

These derived events no longer need to be ordered, since the ordering information is made explicit within the events. These events have the form of very specific if-then clauses. This facilitates the model-fitting process in Ring 1.

(B) For the DNF model with lookback parameter  $L$ , the sequence derived in Ring 3 is blocked in a very similar manner, except that only the selectors describing  $q_i$  are retained in the description of block  $b_i$ . The derived events have the following form:

$$F_0 \& d(F_0, F_1) \& \dots \& d(F_0, F_L) \& s(F_0, F_1) \& \dots \& s(F_0, F_L).$$

These events are very specific conjuncts that are passed to the  $A^q$  algorithm [21, 22] in Ring 1, where they are generalized to form a DNF description.

(C) For the *periodic* model with period length  $P$  and lookback  $L$ , Ring 2 performs a splitting transformation followed by a blocking transformation. First, the sequence obtained from Ring 3 is split into  $P$  separate sequences.

Derived event	Abbreviation
$[\text{rank1}(q_i - 1) = 4][\text{suit1}(q_i - 1) = D]$	$F_1$
$[\text{rank2}(q_i - 2) = 2][\text{suit2}(q_i - 2) = H]$	$F_2$
$\Rightarrow$	
$[\text{rank0}(q_i) = 6][\text{suit0}(q_i) = S]$	$F_0$
$[\text{d-rank01}(q_i, q_i - 1) = 2][\text{d-suit}(q_i, q_i - 1) = 2]$	$d(F_0, F_1)$
$[\text{d-rank02}(q_i, q_i - 2) = 4][\text{d-suit}(q_i, q_i - 2) = 1]$	$d(F_0, F_2)$
$[\text{s-rank01}(q_i, q_i - 1) = 10]$	$s(F_0, F_1)$
$[\text{s-rank02}(q_i, q_i - 2) = 8]$	$s(F_0, F_2)$
$[\text{rank1}(q_i - 1) = 6][\text{suit1}(q_i - 1) = S]$	$F_1$
$[\text{rank2}(q_i - 2) = 4][\text{suit2}(q_i - 2) = D]$	$F_2$
$\Rightarrow$	
$[\text{rank0}(q_i) = 8][\text{suit0}(q_i) = C]$	$F_0$
$[\text{d-rank01}(q_i, q_i - 1) = 2][\text{d-suit}(q_i, q_i - 1) = 1]$	$d(F_0, F_1)$
$[\text{d-rank02}(q_i, q_i - 2) = 4][\text{d-suit}(q_i, q_i - 2) = 3]$	$d(F_0, F_2)$
$[\text{s-rank01}(q_i, q_i - 1) = 14]$	$s(F_0, F_1)$
$[\text{s-rank02}(q_i, q_i - 2) = 12]$	$s(F_0, F_2)$

FIG. 11. Sample events showing which sum and difference descriptors are derived.

Then each separate sequence is blocked into blocks of length  $L + 1$ . The derived events have the same form as the events derived for the DNF model. Note that because the blocking occurs after the splitting, the lookback takes place only within a phase.

To provide an example of the function of Ring 2, Fig. 12 shows some events from Fig. 9 after they have been transformed in preparation for fitting to a decomposition model with  $L = 1$ .

### 6.3.5. Ring 1: the basic model-fitting algorithms

Ring 1 consists of three separate model-fitting algorithms: the  $A^q$  algorithm, the decomposition algorithm and the periodic algorithm.

The  $A^q$  algorithm [21, 22] is applied to fit the DNF model to the data.  $A^q$  attempts to find the DNF description with the fewest number of conjunctive terms that covers all of the positive examples and none of the negative examples. The algorithm operates as follows. First, a positive example, called the seed, is chosen, and the set of maximally-general conjunctive expressions consistent with this seed and all of the negative examples is computed. This set is called a *star*, and it is equivalent to the G-set in Mitchell's [26] version space approach (if the G-set is computed with the seed positive example and all of the negative examples). One element from this star is chosen to be a conjunct in the output DNF description, and all positive examples covered by it are removed from further consideration. If any positive examples remain, the process is repeated, selecting as a new seed some positive example that was not

Derived event	Positive or negative
$\begin{aligned} &[\text{rank1}(b2) = 3][\text{suit1}(b2) = H] \\ &[\text{parity1}(b2) = \text{odd}][\text{color1}(b2) = \text{red}] \\ &[\text{prime1}(b2) = Y][\text{faced1}(b2) = N] \\ &\quad \Rightarrow \\ &[\text{rank0}(b2) = 9][\text{suit0}(b2) = ][\text{parity0}(b2) = \text{odd}] \\ &[\text{color0}(b2) = \text{black}][\text{prime0} = N] \\ &[\text{faced0}(b2) = N][\text{d-rank01}(b2) = +6] \\ &[\text{d-suit01}(b2) = +1][\text{d-parity01}(b2) = N] \\ &[\text{d-color01}(b2) = Y][\text{s-prime01}(b2) = Y] \\ &[\text{d-faced01}(b2) = Y][\text{s-rank01}(b2) = 12] \end{aligned}$	+
$\begin{aligned} &[\text{rank1}(b3) = 9][\text{suit1}(b3) = S] \\ &[\text{parity1}(b3) = \text{odd}][\text{color1}(b3) = \text{black}] \\ &[\text{prime1}(b3) = N][\text{faced1}(b3) = N] \\ &\quad \Rightarrow \\ &[\text{rank0}(b3) = J][\text{suit0}(b3) = D][\text{parity0}(b3) = \text{odd}] \\ &[\text{color0}(b3) = \text{red}][\text{prime0}(b3) = Y] \\ &[\text{faced0}(b3) = Y][\text{d-rank01}(b3) = +2] \\ &[\text{d-suit01}(b3) = +2][\text{d-parity01}(b3) = N] \\ &[\text{d-color01}(b3) = Y][\text{d-prime01}(b3) = Y] \\ &[\text{d-faced01}(b3) = Y][\text{s-rank01}(b3) = 20] \end{aligned}$	-

FIG. 12. Some events of Fig. 9 transformed for decomposition  $L = 1$ .

covered by *any* member of any preceding star. In this manner, a DNF description with few conjunctive terms is found. If the stars are computed without any pruning, then  $A^q$  can provide a tight bound on the number of conjuncts that would appear in the shortest DNF description (i.e., with fewest conjunctive terms).

The decomposition algorithm is an iterative algorithm that seeks to fit the data to a decomposition model. The key task of the decomposition algorithm is to identify a few attributes, called *decomposition attributes*, from which the decomposition rule can be developed. A *decomposition attribute* is an attribute that appears on the left-hand side of an if-then clause of a decomposition rule. For example, the decomposition rule

$$\begin{aligned} &[\text{color}(\text{cardi} - 1) = \text{black}] \Rightarrow [\text{parity}(\text{cardi}) = \text{odd}] \ \& \\ &[\text{color}(\text{cardi} - 1) = \text{red}] \Rightarrow [\text{parity}(\text{cardi}) = \text{even}] \end{aligned}$$

decomposes on COLOR. Hence, COLOR is the single decomposition attribute.

The algorithm uses a generate-and-test approach of the following form:

$$\text{decompositionattributes} := \{ \} \quad \textit{The empty set}$$

```

while rule is not consistent do
  begin
    Generate a trial decomposition
    (based on positive evidence only)
    for each possible decomposition attribute.

    Test these trial decompositions against
    the data.

    Select the best decomposition attribute and
    add it to the set decompositionattributes.
  end

```

The process of generating a trial decomposition takes place in two steps. First, a VL22 conjunction is formed for each possible value of the decomposition attribute. All positive events that have the same value of the decomposition attribute on their left-hand sides are merged together to form a single conjunction of selectors. This VL22 conjunction forms the right-hand side of a single clause in the decomposition rule. Within this conjunction, a selector is created for each attribute by forming the internal disjunction of the values in the corresponding selectors in the events. For example, using all of the events derived in Ring 2 for the sample layout in Fig. 8, the decomposition algorithm generates the trial decomposition shown in Fig. 13 for the  $\text{PARITY}(\text{card}_{i-1})$  attribute.

Since there are only two values (ODD and EVEN) for the decomposition attribute in the sequence shown in Fig. 8, two conjunctions are formed. The first conjunction is obtained by merging all of the positive events

Since there are only two values (ODD and EVEN) for the decomposition attribute in the sequence shown in Fig. 8, two conjunctions are formed. The first conjunction is obtained by merging all of the positive events for which

$$[\text{parity}(\text{card}_i - 1) = \text{odd}] .$$

There are four such events. The first selector in that conjunction,

$$[\text{rank}(\text{card}_i) = 9 \vee 4 \vee 2] ,$$

is obtained by forming the internal disjunction of the values of  $\text{RANK}(\text{card}_i)$  in each of the four events.

The second step in forming a trial decomposition is to generalize each clause in the trial rule. The generalization is accomplished by applying rules of

```

[parity(cardi - 1) = odd]
  ⇒
[rank(cardi) = 9 ∨ 4 ∨ 2]
[suit(cardi) = S ∨ C][parity(cardi) = even ∨ odd]
[color(cardi) = black][prime(cardi) = Y ∨ N]
[faced(cardi) = N]
[d-rank(cardi, cardi - 1) = +6 ∨ -5 ∨ -7]
[d-suit(cardi, cardi - 1) = 1 ∨ 2 ∨ 3]
[d-parity(cardi, cardi - 1) = Y ∨ N]
[d-color(cardi, cardi - 1) = Y ∨ N]
[d-prime(cardi, cardi - 1) = Y ∨ N]
[d-faced(cardi, cardi - 1) = Y ∨ N]
[s-rank(cardi, cardi - 1) = 12 ∨ 13 ∨ 9] &

[parity(cardi - 1) = even]
  ⇒
[rank(cardi) = J ∨ 10 ∨ 8 ∨ 7]
[suit(cardi) = H ∨ D][parity(cardi) = even ∨ odd]
[color(cardi) = red][prime(cardi) = Y ∨ N]
[faced(cardi) = Y ∨ N]
[d-rank(cardi, cardi - 1) = 7 ∨ 8 ∨ -2 ∨ -1]
[d-suit(cardi, cardi - 1) = 0 ∨ 1]
[d-parity(cardi, cardi - 1) = Y ∨ N]
[d-color(cardi, cardi - 1) = Y ∨ N]
[d-prime(cardi, cardi - 1) = Y ∨ N]
[d-faced(cardi, cardi - 1) = Y ∨ N]
[s-rank(cardi, cardi - 1) = 15 ∨ 12 ∨ 18]

```

FIG. 13. Trial decomposition on the PARITY(card<sub>i-1</sub>) attribute.

generalization to extend internal disjunctions and drop selectors. (See [23] for a description of various rules of generalization.) Corresponding attributes in the different clauses of the decomposition rules are compared, and selectors whose value sets overlap are dropped. When these rules of generalization are applied to the trial decomposition of, for example, PARITY, the following generalized trial decomposition is obtained:

$$\begin{aligned}
 &[\text{parity}(\text{cardi} - 1) = \text{odd}] \Rightarrow [\text{suit}(\text{cardi}) = C \vee S][\text{color}(\text{cardi}) = \text{black}] \& \\
 &[\text{parity}(\text{cardi} - 1) = \text{even}] \Rightarrow [\text{suit}(\text{cardi}) = H \vee D][\text{color}(\text{cardi}) = \text{red}].
 \end{aligned}$$

This is a very promising trial decomposition. However, it has been developed using only positive evidence—without considering the possibility that it may cover some negative events. Hence, the trial decomposition must be tested against the negative events to determine whether or not it is consistent. It turns out that the generalized trial decomposition shown above is indeed consistent with the negative evidence.

After a trial decomposition has been developed for each possible decomposition attribute, the best decomposition attribute is selected according to a

heuristic attribute-quality functional. The attribute-quality functional tests such things as the number of negative events covered by the trial decomposition, the number of clauses with non-null right-hand sides, and the complexity of the trial decomposition (defined as the number of selectors that cannot be written with a single operator and a single value). The chosen trial decomposition forms a candidate sequence-prediction rule.

If the candidate rule is not consistent with the data (i.e., still covers some negative examples), then the decomposition algorithm must be repeated to select a second attribute to add to the left-hand sides of the if-then clauses. This has the effect of splitting each of the if-then clauses into several more if-then clauses. For example, if we first decomposed on  $\text{PARITY}(\text{card}_{i-1})$  and then on  $\text{FACED}(\text{card}_{i-1})$ , we would obtain four if-then clauses of the form:

$$\begin{aligned} &[\text{parity}(\text{card}_i - 1) = \text{odd}][\text{faced}(\text{card}_i - 1) = \text{N}] \Rightarrow \dots \\ &[\text{parity}(\text{card}_i - 1) = \text{odd}][\text{faced}(\text{card}_i - 1) = \text{Y}] \Rightarrow \dots \\ &[\text{parity}(\text{card}_i - 1) = \text{even}][\text{faced}(\text{card}_i - 1) = \text{N}] \Rightarrow \dots \\ &[\text{parity}(\text{card}_i - 1) = \text{even}][\text{faced}(\text{card}_i - 1) = \text{Y}] \Rightarrow \dots \end{aligned}$$

The periodic algorithm is similar to the decomposition algorithm. For each phase of the period, it takes all of the positive events in that phase and combines them to form a single conjunct by forming the internal disjunction of all of the value sets of corresponding selectors. Next, rules of generalization are applied to extend internal disjunctions and drop selectors. Finally, corresponding attributes in different phases are compared, and selectors whose value sets overlap are dropped if this can be done without covering any negative examples.

### 6.3.6. *Evaluating the NDP rules*

Once Ring 1 has instantiated the parameterized models to produce a set of rules, the rules are passed back through the concentric rings of the program. Each ring evaluates the rules according to plausibility criteria based on knowledge available in that ring. Ring 2, for example, applies knowledge of the fact that valid sequences can be continued indefinitely. It checks to see that the rule predicts that the sequence could be so-continued. Ring 3—which applies the segmentation transformation—applies its knowledge about the tail end of the unsegmented sequence to make sure it is consistent with the rule. (Recall that the segmentation transformation is sometimes unable to segment the last few events in the sequence.) Ring 4 tests the rule using the plausibility criteria for Eleusis. These criteria are:

(1) Prefer rules with intermediate degree of complexity. In Eleusis, Occam's Razor does not always apply. The dealer is unlikely to choose a rule that is extremely simple, because it would be too easy to discover. Very complex rules will not be discovered by anyone, and, since the rules of the game discourage such an outcome, the dealer is not likely to choose such complex rules either.

(2) Prefer rules with an intermediate degree of nondeterminism. Rules with a low degree of nondeterminism lead to many incorrect plays, thus rendering them easy to discover. Rules that are very nondeterministic generally lead to few incorrect plays and are therefore difficult to discover.

Rules that do not satisfy these heuristic criteria are discarded. The remaining rules are returned to Ring 5 where they are printed for the user.

## 7. Examples of Program Execution

In this section, we present some example Eleusis games and the corresponding sequence-generating laws that were discovered by SPARC/E. Each of these games was an actual game among people, and the rules are presented as they were displayed by SPARC/E (with minor typesetting changes).

The raw sequences presented to SPARC/E had only two attributes: SUIT and RANK. SPARC/E was given definitions of the following derivable attributes:

- COLOR (red for Hearts and Diamonds; black for Clubs and Spades);
- FACE (true if card is a faced, picture card, false otherwise);
- PRIME (true if card has a prime rank, false otherwise);
- MOD2 (the parity value of the card, 0 if card is even, 1 otherwise);
- MOD3 (the rank of the card modulo 3);
- LENMOD2 (when SPARC/E segments the main sequence into derived sub-sequences, it computes the LENGTH of each of the subsequences modulo 2).

Three examples of the program execution are presented. Here are some points to notice in reading the examples. First, each rule is assumed to be universally quantified over all events in the sequence. This quantification is not explicitly printed. Second, when the value set of a selector includes a set of adjacent values, e.g. [rank(card) = 3  $\vee$  4  $\vee$  5], this is printed as [RANK(CARD) = 3..5]. The computation times given are for an implementation in PASCAL on the CDC CYBER 175.

### 7.1. Example 1

In this example, we show the program discovering a segmented rule. The program was presented with following layout:

```

Main line:  AH  7C  6C  9S  10H  7H  10D  JC  AD
Side lines:           KD           5S      QD
                   JH

continued:  4H  8D  7C  9S  10C  KS  2C  10S  JS
           3S           9H      QH
                   6H      AD

```

The program only discovered one rule for this layout, precisely the rule that the dealer had in mind (1.2 seconds required):



RULE 1: LOOKBACK:0 NPHASES:1 PERIODIC MODEL  
 SEGMENTATION CONDITION = [COLOR(CARDI) = COLOR(CARDI - 1)]:  
 PERIOD([LENMOD2(PH1SEGMENTI) = 1])

The rule states that one must play strings of cards with the same color. The strings must always have odd length. The segmentation condition states that a segment is a string of cards all of the same color. CARDI refers to the *I*th card in the original sequence. SPARC/E discovered this rule as a degenerate periodic rule with a period length, *P*, of 1. Hence, PH1SEGMENTI refers to the *I*th segment in phase 1 (the only phase) of the derived sequence. Actually, the rule that the dealer had in mind had one additional constraint: a queen must not be played adjacent to a jack or king.

## 7.2. Example 2

The second example requires the program to discover a fairly simple periodic rule. Here is the layout:

Main line:	JC	4D	QH	3S	QD	9H	QC	7H	QD	
Side lines:	KC	5S				4S		10D		
		7S								
Main line continued:		9D	QC	3H	KH	4C	KD	6C	JD	8D
Main line continued:	JH	7C	JD	7H	JH	6H	KD			

The program discovered three equivalent versions of the rule, which can be paraphrased as "Play alternating faced and nonfaced cards". Here are the rules (0.49 seconds were required):

RULE 1: LOOKBACK: 1 NPHASES: 0 DECOMPOSITION MODEL

[FACE(CARDI - 1) = FALSE] ⇒  
 [RANK(CARDI) ≥ JACK]  
 [RANK(CARDI) > RANK(CARDI - 1)]  
 [FACE(CARDI) = TRUE] &  
 [FACE(CARDI - 1) = TRUE] ⇒  
 [RANK(CARDI) = 3..9]  
 [RANK(CARDI) < RANK(CARDI - 1)]  
 [FACE(CARDI) = FALSE]

RULE 2: LOOKBACK: 1 NPHASES: 1 PERIODIC MODEL

PERIOD([RANK(PH1CARDI) ≥ 3]  
 [RANK(PH1CARDI) ≠ RANK(PH1CARDI - 1)]  
 [FACE(PH1CARDI) ≠ FACE(PH1CARDI - 1)])

RULE 3: LOOKBACK: 1 NPHASES: 2 PERIODIC MODEL

PERIOD([RANK(PH1CARDI) ≥ JACK]  
 [RANK(PH1CARDI) ≥ -RANK(PH1CARDI - 1) + 20]  
 [FACE(PH1CARDI) = TRUE],  
 [RANK(PH2CARDI) = 3..9]  
 [RANK(PH2CARDI) = -RANK(PH2CARDI - 1) + 5..14]  
 [FACE(PH2CARDI) = FALSE])

Rule 1 is a decomposition rule with a lookback of 1. Rule 2 expresses the rule as a degenerate periodic rule with a single phase. Rule 3 expresses the rule in the 'natural' way as a periodic rule of length 2.

Notice that, although the program has the gist of the rule, it has discovered a number of redundant conditions. For example, in Rule 1, the program did not use knowledge of the fact that  $[\text{rank}(\text{cardi}) \geq \text{jack}]$  implies  $[\text{face}(\text{cardi}) = \text{true}]$ , and therefore, it did not remove the former selector. Similarly, because of the interaction of the two conditions in Rule 1,

$$[\text{rank}(\text{cardi}) > \text{rank}(\text{cardi} - 1)]$$

is completely redundant.

### 7.3. Example 3

The third example shows the upper limits of the program's abilities. During this game, only one of the human players even got close to guessing the rule, yet the program discovers a good approximation of the rule using only a portion of the layout that was available to the human players. Here is the layout:

Main line:	4H	5D	8C	JS	2C	5S	AC	5S	10H
Side lines:	7C	6S	KC	AH		6C		AS	
	JH	7H	3H	KD					
	4C	2C		QS					
	10S	7S							
	8H	6D							
	AD	6H							
	2D	4C							

The program produced the following rules after 6.5 seconds:

RULE 1: LOOKBACK: 1 NPHASES: 0 DNF MODEL

$[\text{RANK}(\text{CARDI}) \leq 5][\text{SUIT}(\text{CARDI}) = \text{SUIT}(\text{CARDI} - 1) + 1] \vee$

$[\text{RANK}(\text{CARDI}) \geq 5][\text{SUIT}(\text{CARDI}) = \text{SUIT}(\text{CARDI} - 1) + 3]$

RULE 2: LOOKBACK: 1 NPHASES: 1 PERIODIC MODEL

PERIOD( $[\text{RANK}(\text{PH1CARDI}) = \text{RANK}(\text{PH1CARDI} - 1) - 9]$

$[\text{RANK}(\text{PH1CARDI}) = -\text{RANK}(\text{PH1CARDI} - 1) + 4, 5, 7, 11, 13, 17]$

$[\text{SUIT}(\text{PH1CARDI}) = \text{SUIT}(\text{PH1CARDI} - 1) + 1, 2, 3]$ )

RULE 3: LOOKBACK: 1 NPHASES: 2 PERIODIC MODEL

PERIOD( $[\text{RANK}(\text{PH1CARDI}) = \text{ACE}, 2, 8, 10]$

$[\text{RANK}(\text{PH1CARDI}) = -\text{RANK}(\text{PH1CARDI} - 1) + 1, 8, 9, 10]$ ,

$[\text{RANK}(\text{PH2CARDI}) = 5 \dots \text{JACK}][\text{SUIT}(\text{PH2CARDI}) = \text{SPADES}]$

$[\text{RANK}(\text{PH2CARDI}) = \text{RANK}(\text{PH2CARDI} - 1) + -0 \dots 6]$

```

[RANK(PH2CARDI) = -RANK(PH2CARDI - 1) + 8 .. 14]
[SUIT(PH2CARDI) = SUIT(PH2CARDI - 1) + 0 .. 2]
[COLOR(PH2CARDI) = BLACK][PRIME(PH2CARDI) = PTRUE]
[PRIME(PH2CARDI) = PRIME(PH2CARDI - 1)]
[MOD2(PH2CARDI) = 1][MOD2(PH2CARDI) = MOD2(PH2CARDI - 1) + 0]
[MOD2(PH2CARDI) = -MOD2(PH2CARDI - 1) + 0][MOD3(PH2CARDI) = 2]
[MOD3(PH2CARDI) = MOD3(PH2CARDI - 1) + 0]
[MOD3(PH2CARDI) = -MOD3(PH2CARDI - 1) + 1]

```

The rule that the dealer had in mind was:

```

[SUIT(CARDI) = SUIT(CARDI - 1) + 1]
[RANK(CARDI) ≥ RANK(CARDI - 1)] ∨

```

```

[SUIT(CARDI) = SUIT(CARDI - 1) + 3]
[RANK(CARDI) ≤ RANK(CARDI - 1)]

```

There is a strong symmetry in this rule: the players may either play a higher card in the next 'higher' suit (recall that the suits are cyclically ordered) or a lower card in the next 'lower' suit. The program discovered a slightly simpler version of the rule (Rule 1) that happened to be consistent with the training instances. Note that adding 3 to the `SUIT` has the effect of computing the next lower suit.

The other two rules discovered by the program are very poor. They are typical of the kinds of rules that the program discovers when the model does not fit the data very well. Both rules are filled with irrelevant descriptors and values. The current program has very little ability to assess how well a model fits the data. These rules should not be printed by the program since they are highly implausible.

## 8. Summary

We have presented here a methodology for discovering sequence-generating rules for the nondeterministic prediction problem. The main ideas behind this methodology are

- (1) the use of task-oriented transformations of the initial data and
- (2) the use of different rule models to guide the search for sequence-generating rules.

Four different task-oriented transformations (adding attributes, blocking, splitting into phases, and segmenting) and three models (DNF, periodic, and decomposition) have been presented.

The main part of the methodology has been implemented in the program SPARC/E and applied to the NDP problem that arises in the card game Eleusis. The performance of the program indicates that it can discover quite complex and interesting rules.

This methodology is quite general and can be applied to other nondeter-

ministic prediction problems in which the objects in the initial sequence are describable by a small set of finite-valued attributes. The main strengths of the method are (a) that it can solve learning problems in which the initial training instances require substantial task-oriented transformation and (b) that it can search very large spaces of possible rules using a set of rule models for guidance.

Many aspects of this methodology remain to be investigated. We have not considered NDP problems in which (a) the training instances are noisy, (b) the training instances have internal structure so that an attribute vector representation is not adequate, and (c) the sequence-generating rules are permitted to have exceptions. Application of this methodology to real-world problems will probably also require the development of additional sequence transformations and rule models. Also, more heuristics need to be developed that can be used to guide the application of transformations and models.

The implementation of the methodology in program SPARC/E and the experiments made so far have demonstrated that the method can discover the majority of Eleusis secret rules played in ordinary human games. There are some shortcomings of the implementation, however. The program presently conducts a nearly exhaustive depth-first search of the possible models and transformations. Much could be gained by having the program conduct a best-first heuristically-guided search instead. Another weakness of SPARC/E is its poor ability to evaluate the plausibility of the rules it discovers. It is also not able to simplify rules by removing redundant selectors, nor is it able to estimate the degree of nondeterminism of the rule. Both of these can be implemented without too much difficulty by including inference routines that make more complete use of the background knowledge already available to the program. Finally, an important weakness of the current program is its inability to form composite models. SPARC/E is not presently able to handle the NDP problem shown in Fig. 1, because it involves a periodic rule in which one of the phases contains an imbedded periodic sequence (see Section 5.3).

In addition to these specific problems, there are some more general problems that further research in the area of sequence-generating laws should address. First, in some real-world problems, there are several example sequences available for which the sequence-generating law is believed to be the same. Such problems occur, in particular, in describing the process of disease development in medicine and agriculture. A specific problem of this type that has been partially investigated involves predicting the time course of cutworm infestation in a cornfield and estimating the potential damage to the crop (see [3, 4, 7]). In this problem, several sequences of observations are available—one for each field—and there is a need to develop a sequence-generating law that predicts all of these sequences.

A second general problem for further research is to handle continuous processes. AI research has so far given little attention to this case.

### Appendix A. Notational Conventions

The following notational conventions are employed in this paper. In general, lowercase letters denote objects in some sequence ( $q, ph, b$ ) or index variables ( $i, j, k$ ) or the lengths of sequences ( $m, n$ ). Uppercase letters denote sets of objects, attributes, and so on ( $Q, F, S$ ) as well as parameters of models and transformations ( $L, P$ ). Small capitals denote attributes (COLOR, RANK) and their values (RED, KING).

- $\langle \rangle$  Angle brackets denote sequences of objects, e.g.,  $\langle 2, 4, 6, 8 \rangle$  and also periodic rules, e.g.,  $\langle [\text{color}(\text{ph}1i) = \text{red}], [\text{color}(\text{ph}2i) = \text{black}] \rangle$ .
- $q_i$  or  $q_i$  The  $i$ th object in an input sequence.
- $q'_i$  The  $i$ th object in a derived sequence.
- $q_{\bar{i}}$  An object that constitutes an incorrect extension of the sequence after object  $q_{i-1}$ .
- $b_i$  or  $b_i$  The  $i$ th block in a sequence derived by the blocking transformation.
- $ph_i$  The  $i$ th phase derived by the splitting transformation.
- $ph_{i,j}$  or  $ph_{ij}$  The  $j$ th object in the  $i$ th phase after a splitting transformation.
- $n$  The number of descriptors.
- $E$  The space of possible events.
- $F$  The starting set of attributes for a transformation.
- $S$  The starting set of sequences for a transformation.
- $Q$  The starting set of objects for a transformation.
- $F'$  The set of derived attributes from a transformation.
- $S'$  The set of derived sequences from a transformation.
- $Q'$  The set of derived objects from a transformation.
- $F_j$  The set of selectors describing object  $q_{i-j}$  in block  $b_i$ .
- $g$  The sequence-generating function that maps a sequence into a set of objects  $Q_{k+1}$  that can appear as continuations of the sequence.
- $Q_{k+1}$  The set of objects that can appear as continuations of the sequence  $\langle q_1, q_2, \dots, q_k \rangle$ .
- $P$  The number of phases parameter of the splitting transformation and the period model.
- $L$  The lookback parameter of the blocking transformation and all three models.
- $[f_i(q_j) = r_k]$  or  $[f_i(q_j) = r_k]$  A simple selector, which asserts that feature  $f_i$  of object  $q_j$  has the value  $r_k$ .
- $[f_i(q_j) = r_1 \vee r_2 \vee r_3]$  A selector containing an internal disjunction. It asserts that  $f_i$  can have the value  $r_1$  or  $r_2$  or  $r_3$ .
- $d(\text{prefix})$  The  $d$  prefix on an attribute name indicates that it is a

- difference attribute. Hence,  $D-RANK(q_i, q_{i-1})$  is equal to  $RANK(q_i) - RANK(q_{i-1})$ .
- $s$  (prefix) The  $s$  prefix on an attribute name indicates that it is a summation attribute. Hence,  $S-RANK(q_i, q_{i-1})$  is equal to  $RANK(q_i) + RANK(q_{i-1})$ .
- $d(F_i, F_j)$  The set of difference selectors obtained by 'subtracting' selectors  $F_j$  from  $F_i$ .
- $s(F_i, F_j)$  The set of summation selectors obtained by 'adding' selectors  $F_i$  and  $F_j$ .
- $\Rightarrow$  Logical implication.

#### ACKNOWLEDGMENT

The authors wish to thank the referees of the AI Journal for their suggestions. Thanks also go to the referees of the 1983 Machine Learning Workshop where an earlier version of this paper appeared. The authors gratefully acknowledge the partial support of the NSF under grant MCS-82-05166 and of the Office of Naval Research under grant No. N00014-82-K-0186.

#### REFERENCES

1. Abbott, R., The new Eleusis, Available from the author, Box 1175, General Post Office, New York, NY 10116, 1977.
2. Amarel, S., On representations of problems of reasoning about actions, in: D. Michie (Ed.), *Machine Intelligence 3* (University of Edinburgh Press, Edinburgh, 1968) 131-171.
3. Baim, P.W., Automated acquisition of decision rules: Problems of attribute construction and selection, M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, 1983.
4. Boulanger, A.G., The expert system PLANT/CD: A case study in applying the general purpose inference system ADVISE to predicting black cutworm damage in corn, M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, 1983.
5. Buchanan, B.G. and Mitchell, T.M., Model-directed learning of production rules, in: D.A. Waterman and F. Hayes-Roth (Eds.), *Pattern-directed Inference Systems* (Academic Press, New York, 1978).
6. Cohen, P. and Feigenbaum, E.A., *The Handbook of Artificial Intelligence 3* (Kaufmann, Los Altos, CA, 1982).
7. Davis, J., CONVART: A program for constructive induction on time-dependent data, M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, 1981.
8. Dietterich, T.G., London, R., Clarkson, K. and Dromey, G., Learning and inductive inference, in: P.R. Cohen and E.A. Feigenbaum (Eds.), *The Handbook of Artificial Intelligence 3* (Kaufmann, Los Altos, CA, 1982) Ch. 14.
9. Dietterich, T.G. and Michalski, R.S., Inductive learning of structural descriptions: Evaluation criteria and comparative review of selected methods, *Artificial Intelligence* 16 (1981) 257-294.
10. Englemore, R. and Terry, A., Structure and function of the crystal system, *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, 1979.

11. Friedland, P.E., Knowledge-based experiment design in molecular genetics, Rept. No. HPP-79-29, Department of Computer Science, Stanford University, Stanford, CA, 1979.
12. Gardner, M., On playing the new Eleusis, the game that simulates the search for truth, *Scientific American* 237 (1977) 18-25.
13. Hedrick, C.L., Learning production systems from examples, *Artificial Intelligence* 7 (1976) 21-49.
14. Karpinski, J. and Michalski, R.S., A system that learns to recognize hand-written alphanumeric characters, *Proc Institute Automatyki, Polish Academy of Sciences* 35 (1966).
15. Kotovsky, K. and Simon, H.A., Empirical tests of a theory of human acquisition of concepts for sequential patterns, *Cognitive Psychol.* 4 (1973) 399-424.
16. Langley, P.W., Descriptive discovery processes: experiments in Baconian science, Rept. No. CMU-CS-80-121, Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA, 1980.
17. Lenat, D.B., The nature of heuristics, in: R.S. Michalski, J.G. Carbonell and T.M. Mitchell (Eds.), *Machine Learning* (Tioga, Palo Alto, CA, 1983).
18. McCarthy, J., Programs with common sense, in: *Proceedings of the Symposium on the Mechanization of Thought Processes*, National Physical Laboratory, (1958) 77-84.
19. McCarthy, J., Epistemological problems of artificial intelligence, *Proceedings of the Fifth International Joint Conference on Artificial Intelligence* (1977) 1038-1044.
20. McCarthy, J. and Hayes, P., Some epistemological problems from the standpoint of artificial intelligence, in: B. Meltzer and D. Michie (Eds.), *Machine Intelligence* 4 (Edinburgh University Press, Edinburgh, 1969) 463-502.
21. Michalski, R.S., On the quasi-minimal solution of the general covering problem, in: *V International Symposium on Information Processing, FCIP*. 69, Yugoslavia, Vol. A3 (1969) 2-12.
22. Michalski, R.S., A variable-valued logic system as applied to picture description, in: P. Nake and A. Rosenfeld (Eds.), *Graphic Languages* (North-Holland, Amsterdam, 1972) 20-47.
23. Michalski, R.S., A theory and methodology of inductive learning, *Artificial Intelligence* 20 (1983) 111-161.
24. Michalski, R.S., Carbonell, J.G. and Mitchell, T.M. (Eds.), *Machine Learning* (Tioga, Palo Alto, CA, 1983).
25. Michalski, R.S. and Chilausky, R.L., Learning by being told and learning from examples: An experimental comparison of the two methods of knowledge acquisition in the context of developing an expert system for soybean disease diagnosis, *Policy Analysis and Information Systems* 4 (2) (1980).
26. Mitchell, T.M., Version spaces: An approach to concept learning, Rept. No. STAN-CS-78-711, Stanford, CA, 1978.
27. Mitchell, T.M., Utgoff, P.E. and Banerji, R.B., Learning by experimentation: Acquiring and refining problem-solving heuristics, in: R.S. Michalski, J.G. Carbonell and T.M. Mitchell (Eds.), *Machine Learning* (Tioga, Palo Alto, CA, 1983).
28. Persson, S., Some sequence extrapolating programs: A study of representation and modeling in inquiring systems, Rept. No. STAN-CS-66-050, Department of Computer Science, Stanford University, Stanford, CA, 1966.
29. Samuel, A.L., Some studies in machine learning using the game of checkers, in: E.A. Feigenbaum and J. Feldman (Eds.), *Computers and Thought* (McGraw-Hill, New York, (1963) 71-105.
30. Samuel, A.L., Some studies in machine learning using the game of checkers II—Recent progress, *IBM J. Res. Development* 11 (6) (1967) 601-617.
31. Schank, R. and Abelson, R., Scripts, plans, and knowledge, *Proceedings of the Fourth International Joint Conference on Artificial Intelligence* (1975) 151-157.
32. Simon, H.A., Complexity and the representation of patterned sequences of symbols, *Psych. Rev.* 79 (5) (1972) 369-382.

33. Simon, H.A. and Kotovsky, K., Human acquisition of concepts for sequential patterns, *Psychol. Rev.* 70 (1963) 534-546.
34. Solomonoff, R.S., A formal theory of inductive inference, *Inform. and Control* 7 (1964) 1-22; 224-254.
35. Soloway, E.M., Learning = Interpretation + Generalization: A case study in knowledge-directed learning," Rept. No. COINS TR-78-13, Computer and Information Science Department, University of Massachusetts at Amherst, 1978.
36. Winston, P.H., Learning structural descriptions from examples, Rept. No. AI-TR-231, MIT, Cambridge, MA, 1970.

*Received July 1983; revised version received May 1984*