Order Number 9224829

Implicit detection of event interdependencies

Kushner, Harry Michael, Ph.D.

City University of New York, 1992

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106

IMPLICIT DETECTION OF EVENT INTERDEPENDENCIES

by

HARRY MICHAEL KUSHNER

A dissertation submitted to the Graduate Faculty in Psychology in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

1992

This manuscript has been read and accepted for the Graduate Faculty in Psychology in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

April 30, 1992
Date

Arthur Reber
Chair of Examining Committee

April 30, 1992
Date

Executive Officer

Eric G. Heinemann

David Owen

Donald L. Scarborough

Joan Gay Snodgrass

Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

IMPLICIT DETECTION OF EVENT INTERDEPENDENCIES
by
Harry Michael Kushner

Adviser: Professor Arthur S. Reber

This dissertation presents three experiments which explore the ability of human subjects to extract covariation information from rule-governed visual displays. The procedure used in all of the experiments involves a prediction task in which subjects "guess" (with feedback) the location of a target event following a sequence of 5 or 9 similar events. In all cases, the location of the target event can be predicted with perfect accuracy according to a deterministic rule based on the relationship of 2 of the events in the immediately preceding sequence. Results from the experiments support the following broad conclusions: first, subjects can learn rules of this kind as measured by their steadily improving accuracy; second, their learning is largely implicit, that is, unaccompanied by conscious, explicit knowledge of the rules which govern the correct responses; third, they can transfer this learning to a new rule structure when the feedback is systematically shifted; and fourth, instructions which inform some subjects about the existence of a rule result in larger between-subject variation in performance relative to those not so informed. Implications of this work for implicit learning theory, and for cognitive psychology in general, are discussed.

TABLE OF CONTENTS

List of Tables

List of Tables (continued)

List of Tables (continued)

List of Figures

List of Figures (continued)

INTRODUCTION

Twenty-five years of research has established that subjects confronted with particular kinds of stimulus environments and given appropriate instructions acquire implicit knowledge about the regularities which such environs contain. Further, the process of acquisition itself occurs without the instruction or, apparently, the intention to learn. As has been argued elsewhere (Cleeremans & McClelland, 1991; Reber, 1989), the experiments in this corpus of research all have in common the property that the underlying knowledge base that subjects extract from their interactions with the stimulus environments can be captured by the notion of covariation. That is, subjects' behavior in all of these experiments appears to reflect a single process: the detection of covariation among events as they are instantiated in the stimulus display. The generality of this process is considerable and it has been observed across a wide range of stimulus materials (Lewicki, 1986; Reber & Millward, 1968, 1971).

This dissertation explores this notion of the detection of covariation by introducing a stimulus environment that is based on a complex array of events whose underlying structure is characterized by a remote, deterministic, double-dependency rule that is more complex than anything that has been studied to date. If implicit learning is as robust a process as some have suggested (Lewicki, 1986; Lewicki & Hill, 1989; Reber, 1989), then we ought to be able to observe the process emerging in such situations where the associative links between events are complex and non-salient.

The procedure used in all of the new studies presented here is a relatively simple prediction experiment in which subjects had to "guess" the successor of a sequence of similar events. Subjects were exposed to series of either 5 or 9 stimuli presented successively on a computer monitor and were asked to predict the location of the next (target) stimulus. The location of the target stimulus was determined on the basis of the relationship between the locations of exactly two of the events in the immediately preceding series. The other events in that series were irrelevant for prediction accuracy. There are at least two reasons why this task may be quite hard. First, there are more irrelevant events than useful ones. Second, the rule that defines the location of the target stimulus is complex in that it involves a relationship between events rather than the particular events themselves. This results in each component of the rule being instantiated by different pairs of events, each of which may in turn be embedded in a large number of different irrelevant contexts.

Analysis of the three experiments presented supports the following broad conclusions: first, subjects can indeed learn contexted, double dependencies of this kind, and they do so largely independently of any explicit knowledge of the rules. Second, once the pattern of interdependencies has been detected, subjects are capable of transferring their knowledge to a "shifted" rule. Third, performance dissociations are observed between subjects informed about the existence of a rule in the stimulus display and those not so informed.

CHAPTER 1

Implicit learning is the process whereby knowledge about complex, rule-governed stimulus environments is acquired without specific intentions to learn and largely independently of conscious knowledge about what was learned (Reber, 1989). This process has been explored in a wide variety of experimental contexts including artificial grammar learning (Reber, 1967, 1989), patterned sequence learning (Cleeremans & McClelland, 1991; Lewicki, Hill, & Bizot, 1988; Nissen & Bullemer, 1987), concept formation (Brooks, 1978), probability learning (Millward & Reber, 1972; Reber & Millward, 1968, 1971), and process control of simulated manufacturing plants (Berry & Broadbent, 1984). In all of these cases subjects learn to make decisions, classify novel stimuli, anticipate events, and solve problems that require knowledge of the regularities in the stimulus environment while showing little or no explicit, reportable knowledge about these regularities. Subjects' prediction accuracy or decrease in reaction time in these experiments is significantly above the level expected by chance alone and sometimes dramatically so. Yet these same subjects have little explicit, reportable knowledge of the environments about which they have procedural and functional knowledge. That is, they make accurate judgments about exemplars and events drawn from highly complex, rule-governed stimulus environments but can say little about the rules, regularities, or covariations therein--hence the term "implicit learning". The term was specifically chosen to differentiate this type of learning from the more traditional, overt, hypothesis-testing type of human learning that experimental psychologists typically study (see Reber, 1989 for a review).

The studies noted above share a number of attributes which have been isolated as important criteria for eliciting implicit learning. The first is that regardless of the task requirements, subjects operate within stimulus environments whose events or symbols are governed by complex rule structures--too difficult to be learned explicitly in a typical laboratory session. The second is that, by virtue of the first reason, they all invoke an alternative mode of information processing which "takes over" tasks with such rule structures and allows the subjects to exploit the regularities and covariations therein. What follows is an overview of the research that lends convergent support to these suggestions.

I. Historical and Contemporary Context of the Issue.

An important early study in the area of implicit learning explored the sensitivity of subjects to a rule based on a single remote invariance in probability learning experiments (Millward & Reber, 1972). Subjects observed sequences in which simple events (one or another of two lights being illuminated) were dependent on some previous event that had occurred either 1, 3, 5, or 7 events earlier. They later predicted, with above-chance accuracy, events whose occurrence was based on information from as many as seven trials "back" from the current trial. But when the same subjects were asked to explicitly recall those earlier events, they performed at chance levels of accuracy for events that were 5 or 7 trials back. This finding indicates that the subjects were sensitive to information contained in earlier trials although they had little or no explicit knowledge about those trials.

Recently, further investigation of this ability has explored the possibility that subjects are implicitly sensitive to rules which contain <u>two</u> events with an invariant relation (or "double-dependency" rules). This work employs an extension of the Millward & Reber (1972) methodology. These studies have adapted that methodology by making a to-be-predicted event dependent on two prior events instead of one. The question of interest is whether subjects are sensitive to this more complex manipulation when they are required to predict future events.

This possibility was first explored in a somewhat different context by Mathews, Buss, Stanley, Blanchard-Fields, Cho & Druhan (1989) in an artificial grammar experiment. Artificial grammar experiments, initiated by Reber (1967), represent historically the largest single methodological technique for studying implicit learning. They differ from probability learning experiments in that they use synthetic grammars for the underlying patterns that make up the stimulus environment. The grammars are Markovian systems that generate strings of letters (or any symbols) in a left-to-right fashion (see Figure 1 for an example). The strings represent what would be sentences in a true grammar. In the typical study, subjects, unaware of the existence of the grammar, are exposed to "legal" strings. They acquire knowledge of its rules via some required task which usually involves memorizing exemplars of legal strings to a criterion of, for example, two errorless recalls. After this acquisition phase, subjects' knowledge of those rules is assessed, usually by requiring them to discriminate between grammatical and ungrammatical exemplars of the grammar.

The Mathews et al. (1989) study was primarily concerned with exploring the effects of instructional manipulation to recruit implicit or explicit learning processes. The authors achieved this by dividing their subjects between two instructional groups: one that encouraged subjects to search for the underlying rules of the grammar, and one that had subjects memorize grammatical exemplars without knowing that they were generated by a grammar. In the experiment that investigated the role of the manipulation in learning a two-event rule, they used a grammar with two special attributes. First, it generated strings that had two sets of four letters separated by a period. Second, the grammar "built in" three two-event rules to each string produced: X goes with T, P with C, and S with V. So, for example, the strings XCSS.TPVV and PSTV.CVXS are valid strings because each member of each conditional pair has its associate reflected in the appropriate position across the period. The authors hypothesized that subjects given an acquisition task that encouraged explicit, overt learning processes would perform better in a subsequent recognition task than those given instructions that recruited exclusively implicit processes. Their logic was that first, two-event rules can be explicitly generated, and second, such rules should be difficult to detect implicitly. They found that, indeed, conditions that recruit exclusively implicit learning processes were not sufficient for learning of two-event rules to occur. In addition to its inclusion of an instructional variable, this study is relevant to the dissertation research because it uses three deterministic rules which each involve two events.

Cleeremans & McClelland (1991), in contrast, present evidence that subjects _are_ implicitly sensitive to such a contingency rule under certain circumstances. They too used a finite-state grammar within a very different methodology to investigate the ability of subjects to detect not deterministic pairs of events, but probabilistic interdependencies built into the stimulus display. Instead of generating strings of letters, their grammar (Figure 2) generated the location at which a spot could appear among six potential locations on a computer monitor. As noted above, there is nothing special about letter strings; the grammars can be used to generate, in a Markovian way, other symbols, locations, or any stimulus events. Subjects were required to press a corresponding key on the keyboard in response to the spot appearing at one of the six positions, which were arrayed horizontally across the monitor. Each spot appearance and subsequent subject response represented one trial. Speed and accuracy were stressed, and stimulus generation consisted of "strings" of spot locations generated end-on-end: when the grammar reached node #0 at the far right, it looped back on itself and started at node #0 on the far left. Subjects ran blocks of 155 trials at a time and were not told of any regularity or structure among the sequences of spot appearances.

An important facet of this study is that certain paths taken by the grammar make some subsequent spot locations more likely than others. For example, consider the grammar in Figure 2. Starting at node #0 on the far left, it can generate the paths XTV, PTV, and QTV, which differ only in their first element. In addition, the letter 'S' is legal only after XTV, and 'T' is legal only after PTV and QTV. Cleeremans and McClelland hypothesized that if subjects are sensitive to an interdependency three items removed from the current trial, their reaction time (RT) to respond to the successor of these three paths should be lower than when there is less predictability or no predictability to the path's successor. This hypothesis was confirmed, and prompted further investigation into the limits to the amount of irrelevant material that could be interposed between two crucial events if learning of the relationship between those events is to occur. Specifically, they found that information about events three elements removed from the current trial could be exploited by subjects. When four or more irrelevant elements were interposed between the crucial events, no learning occurred.

Cleeremans and McClelland reported that in post-experimental interviews, all subjects said that they noticed that short sequences of alternating material occurred frequently; all but one of the 6 subjects noticed specific pairs of positions between which the alternating material was occurring. One subject even reported noticing three positions involving an alternation but could not specify them or elaborate when probed further. When asked directly whether they had attempted to exploit the detected regularities to maximize their performance, the subjects reported that they had done so only for short periods because using explicit strategies diminished their accuracy and increased their reaction times. Cleeremans and McClelland concluded that the subjects "had limited reportable knowledge of the sequential structure of the material, and that they tried not to use what little knowledge they had". The most important difference between this study and the Mathews et al. (1989) study is that a remote, probabilistic rule, embedded in the larger context of a grammar, has replaced deterministic rules governing non-contexted letter pairs.

Another contemporary research group investigating covariation detection is headed by Lewicki (Lewicki, Czyzewska, & Hoffman, 1987; Lewicki, Hill, & Bizot, 1988). The basic methodology of this group departs noticeably from probability learning and grammar-based pattern learning. Lewicki et al. (1987) used a complex rule that specified in which quadrant of a monitor a visual target would appear on a search trial. Within a block of seven consecutive trials, subjects were required to locate a target in a display and press a key corresponding to the target's location among the four quadrants. In the first six trials of a block, the target (a single digit) was the only item on the monitor, so the "search" task was quite simple. On the seventh trial, however, the target was embedded among thirty-five other digits, but its quadrant location was determined by its location in the first, third, fourth, and sixth immediately preceding six "simple" trials (its locations in the second and fifth preceding trials were irrelevant). Results showed that 1) subjects were accurate on 98% or more of the trials; 2) subjects' RT to the target's location decreased throughout the experiment; 3) the decrease in RTs was not due to a practice or facilitation effect; 4) when the rule which governed the target's location was changed, RT immediately increased (subjects were told to stress accuracy over speed); and 5) subjects could not specify any part of the rule that governed the target's location, even when a substantial cash bonus was offered for being able to articulate even a fragment of it. Interestingly, in one of the studies, the subjects were all colleagues of Lewicki in the psychology department at the University of Tulsa and were familiar with his interest in issues of nonconscious cognition. A recent critique of this study by Perruchet, Gallego, & Savy (1990) will be discussed later in this chapter.

In addition to the research programs of Reber, Mathews, Lewicki, and their colleagues, there are two research groups that have pursued related lines of inquiry. Although these groups' experiments do not address the issue of two-event rule detection *per se*, they have clear implications for the fundamental issue of the detection of frequency and covariation and the cognitive nature of such detection processes. One such group is headed by Hasher and her colleagues (Hasher & Chromiak, 1977; Hasher, Goldstein, & Toppino, 1977; Hasher & Zacks, 1979, 1984; Zacks, Hasher, & Sanft, 1982). They have found, in a long series of studies, that the acquisition (or "logging") of frequency data is affected minimally or not at all by the extent of practice at such tasks, competing task demands, and the accuracy of test expectations. This last element is important inasmuch as these experiments typically use simple methodology involving free recall of word lists, and subjects are unaware that their sensitivity to the frequency of items on various lists presented will be tested. The main conclusion of the work of Hasher and her colleagues is that frequency information is one of a small number of stimulus attributes (along with temporal and spatial location information) that are encoded automatically, that is, without intention. Furthermore, neither instructions requiring subjects to attend to frequency in a series of word lists, nor training at frequency-logging tasks, improves performance.

The second group, led by Nissen (Nissen & Bullemer, 1987; Willingham, Nissen, & Bullemer, 1989) has studied both normal and memory-impaired (amnesic and Korsakoff's syndrome) individuals using a sequence learning procedure. Subjects in these studies "shadow" (on the keyboard) repeating end-on-end sequences of spots that appear at one of four locations on a computer monitor. The sequences, however, are not generated by a grammar or any probabilistic structure--they are merely repetitions of a fixed ten-location pattern that loops back on itself continuously for the duration of the experiment. Although it has not been possible to study the covariation detection issue with this design, Nissen's work has produced some findings that have implications for this dissertation research. First, subjects' RTs decrease continuously over the experiment, and if the sequence is suddenly and surreptitiously changed, reaction time immediately increases dramatically. Second, subjects showed procedural learning of the sequence in the absence of explicit declarative knowledge of it: they could not generate the repeating sequence despite their nearly 100 millisecond drop in reaction time over forty repetitions of the sequence. Once again, the issue of detection of covariations or remote rules within a display is not addressed, but a more fundamental ability related to it seems to have been discovered with this research.

II. Summary.

Although all of the foregoing work is relevant to the larger issue of implicit learning, the two research programs that are most germane are those of Lewicki and Cleeremans and their coworkers. The pattern-learning experiments of Lewicki and his colleagues are somewhat similar to those of Cleeremans and his coworkers in that they seem to induce nonconsciously acquired procedural knowledge in the absence of declarative knowledge. That is, their subjects seem to be able to exploit the interdependency among four simple events and the target event, much as Cleeremans & McClelland's (1991) subjects exploit the regularities embedded in the grammar they use. Recent work by Perruchet, Gallego, and Savy (1990) shows, however, that the results of Lewicki et al. (1987) can be accounted for by the relative frequency of a small number of "simple" target locations. They argued that since RTs are highly sensitive to the frequency of events in a sequence, they should be directly related to the infrequency of the last movement of a given target. That is, the less often a particular target movement occurs, the longer the RT should be to respond to it. Following up on this logic, they replicated Lewicki et al. (1987) but counterbalanced the frequency of events to eliminate any effects of event probability on RT. Perruchet et al. (1990) reported that they account for all of the decrease in reaction time with the probabilistic nature of the two initial simple search trials within a block of seven. They conclude that this analysis shows that Lewicki et al.'s (1987) claims about nonconscious processing are unfounded since subjects' improvement in performance is unrelated to implicit knowledge. The work of Perruchet et al. (1990) does not, however, completely undermine the implicit knowledge claim: what the subjects learned may be different from what Lewicki et al. (1987) wanted them to learn, but the way it was learned (implicitly) was as they claimed. If this position is correct, adjustments in stimulus control procedures are needed to preclude the types of artifacts that Perruchet et al. (1990) seem to have found. Deeper considerations of the issue of dissociation between the rules and information that subjects learn, versus the rules and information experimenters build into stimulus displays, will be fully addressed later.

The last point about dissociation between rules "in the experiment" and rules "in the mind" make Lewicki's research compelling in that his arrangement of processing demands and task requirements seems to recruit implicit learning processes. Similar to grammar-learning tasks and those of Cleeremans & McClelland, Lewicki's paradigm uses stimulus generating mechanisms which embed remote interdependencies among events. However instead of "building in" those interdependencies via a Markov process, Lewicki composes rules which make the target locations in crucial trials dependent on its locations in certain immediately preceding simple trials. Both mechanisms are sufficiently complex so that conscious and explicit attempts to learn them will generally be unsuccessful. They thus involve the same style of information processing which, in the grammar tasks, has resulted in implicit learning. Concerning the issue of task characteristics, Lewicki's experiments involve a search task, not a grammaticality judgment (like grammar tasks), a sequence-shadowing task (like those of Nissen), or a prediction task (like probability learning tasks). Given his stimulus materials, the first two tasks are probably not applicable. But it is conceivable that in his experiments a prediction task would also evoke implicit learning processes, since his subjects could have been required to predict the crucial target locations instead of searching for them. Such a requirement would have maintained the spirit of his investigation of covert learning processes and might have countered the claims of Perruchet et al. (1990) about his subjects' response strategies.

The work of Cleeremans & McClelland appears not to suffer from the kind of flaw that Perruchet et al. (1990) seem to have discovered in Lewicki et al.'s (1987) study. The part of their work that addresses the issue of event interdependencies is somewhat simpler in that it requires the subject to "carry over" less information within a trial. Lewicki et al.'s (1987) rule is a test of subjects' sensitivity to contingencies across trials and uses a deterministic rule, while Cleeremans & McClelland's (1991) grammar is by definition a probabilistic rule structure that requires less carry-over of information. Theoretically, a combination of these methodologies could recruit implicit learning processes while avoiding flaws due to stimulus generation. The first such combination investigated involved a Lewicki-type rule structure requiring a prediction response instead of a search task, and that is reported here as Experiment 1. This experiment was completed before the Perruchet et al. (1990) criticism of Lewicki et al. (1987) appeared, and it was later discovered that it too had a flaw that, while not entirely eliminating the significance of the experiment, seriously diminished its interpretability. Experiments 2 and 3 sought to eliminate the flaws in, yet retain successful elements from, Experiment 1.

CHAPTER 2

This chapter presents Experiment 1 and a detailed discussion of its results and implications in light of the research presented in the previous chapter.

Experiment 1

Subjects:

Subjects were 5 Brooklyn College students. They received $25.00 for their participation plus an incentive fee of 1 cent for each correct response.

Apparatus and display:

The experiment was run on an IBM microcomputer. The display consisted of three numbered boxes arrayed horizontally across the monitor and measuring 6.5 cm X 7.5 cm each. A trial consisted of 9 randomly generated successive events and a to-be-predicted target event. Each event consisted of the appearance of a square stimulus (2.5 cm wide) in one of the boxes. The stimulus remained on the screen for 500 ms, disappeared, and reappeared in the same or another box 200 ms later (the interstimulus interval) until it had appeared 9 times. After the ninth event had occurred, subjects were asked to predict in which box the stimulus would appear next; they entered their prediction by typing a number between 1 and 3 on the numeric keypad of the keyboard. The computer then displayed the correct response for 3000 ms. Subjects initiated the next trial by pressing the space bar.

<u>Design and method</u>:

The entire experiment consisted of 3000 trials and was broken down into three phases, the existence of which subjects were unaware. Intentionally vague instructions described the experiment as being concerned with "prediction behavior". Phase I (the "training" phase) consisted of 2000 trials. During this phase, the location of the 10th event could be predicted perfectly based on the relationship between the locations at which the third and seventh stimuli of the current trial had appeared. If these stimuli had appeared in the same box, then the tenth stimulus appeared in Box 1; if the seventh appearance was in a box to the right of the third, the tenth stimulus appeared in Box 2; and if the seventh appeared in a box to the left of the third, the tenth stimulus appeared in Box 3. This rule accounts for all of the 9 possible pairs of locations of the third and seventh events (see Table 1, Column I).

On the first 500 trials of the prediction trials, subjects were also given a memory task. After the feedback to the prediction response disappeared, they were asked to report in which box the spot was located on its <u>nth</u> appearance, where <u>n</u> was a randomly generated integer from 1 to 9; they responded using the same keys but no feedback was given.

Over four weeks, each subject progressed through the experiment in blocks of 50 trials at a time at a schedule and pace convenient for them. At a given sitting, they were instructed to run as many blocks as they felt they could give their full and undivided attention. The record of subjects' responses and their reaction times were written to the same disk that contained the experiment. After every 500 trials, the records were downloaded and the disk returned to the subject.

The final 1000 trials consisted of 2 surreptitious tests of the subjects' knowledge of the rule which governed the feedback. When the disk was returned to them after the 2000th trial, the rule was changed so that a pattern that had predicted a tenth location in Box 1 now predicted Box 2; one that predicted Box 2 now predicted Box 3; and one that predicted Box 3 now predicted Box 1. In other words, the correct response was "shifted" one box. Subjects were given no new instructions and told to return the disk after 500 trials as usual. This constituted Phase II, the "transfer" phase (see Table 1, Column II).

For the final 500 trials, the rule governing feedback was again changed to one that generated random "correct" locations after subjects made their prediction responses, that is, there was no programmed connection between any of the sequences' events and their feedback. This was Phase III, the "random" phase (see Table 1, Column III).

After completing this last phase, subjects were asked to give detailed reports of what they were doing in the experiment to maximize their correct responses. They were also questioned extensively about their knowledge of the rules that governed the feedback and asked if they could state explicitly any part of those rules from any phase of the experiment. Finally, each subject was debriefed and paid.

Results:

Figure 3 shows the average proportion of correct responses for each session of the experiment; Figures 4 through 8 show individual performance data for each subject. In Figures 3 through 8, a horizontal line is drawn to indicate chance levels of performance.

Table 2 presents the block-by-block accuracy for each subject in each phase of the experiment, and Table 3 presents the average accuracy, by phase, for each subject.

Subjects' prediction accuracy increases over the first 2000 trials of training (Phase I) and ends up reaching an average of 39.6% correct responses in the final 200 trials of the training phase. Subjects predicting correctly at the level expected by chance would achieve 33.3% accuracy. Within a block of 200 trials, the minimum accuracy necessary for a given subject to achieve significantly above-chance performance is 39.0%; for Phase I as a whole (2000 trials) the minimum figure for a given subject is 35.1% ($p < 0.05$, one-tailed). These figures were calculated using a normal approximation to the binomial distribution. Table 2 shows that all subjects, individually, eventually reach significantly above-chance levels of prediction accuracy and that their average accuracy for Phase I as a whole is significantly above chance. This indicates that subjects have acquired knowledge about the relevant regularities embedded in the material.

The second (shifted rule) phase begins with a dramatic drop in accuracy, but there is some evidence of learning over the course of the phase: subjects eventually reach an average level of 35.1% correct responses in the last 100 trials of Phase II. This figure is <u>not</u> significantly above chance, however ($p > 0.05$; the minimum level necessary for significance is 42.0% in a block of 100 trials). In the 500 trials of Phase II overall, subjects only achieve an average of 33.2% accurate predictions (minimum accuracy necessary for significance is 36.8% in the 500 trials).

As expected, accuracy in the third (random feedback) phase is low relative to accuracy in the last block of Phase II, and although all subjects do exceed chance levels of accuracy (33.3%) in one or more blocks, their individual and overall performance fail to ever reach significantly above-chance levels (42.0% in a block of 500 trials and 36.8% in the 500 trials overall; $p > 0.05$).

Despite the clear sensitivity to the complex regularities embedded in the material in Phase I, none of the subjects exhibited explicit knowledge of the sequential structure relating the crucial events with the feedback. None was able to state the rule in force or any part of the rule during any part of the experiment. For example, when pressed to articulate the exact nature of the strategies they had used in the prediction task, subject #1 reported using a "counting rule", using the box numbers and event numbers as values to be arithmetically combined in order to discover the number of the correct box. He could not, however, articulate the exact nature of the manipulation or whether it involved addition, subtraction, etc., nor the nature of the "values", as he put it, that were assigned to each box. Subject #3 reported assigning each of the ten events to one of the fingers on her hands and holding each finger at a certain height to symbolize the box in which it appeared. In order to discover where the tenth event would appear, she observed her hands and decided what position "looked natural" for the tenth finger given the positions of the other nine. When asked to articulate how she decided what was "natural" about the position she chose for the tenth location, she reported that she discovered a prediction strategy involving events 1, 2, 8, and 9 that was based on her musical training as a pianist; she characterized each prediction task as what finger positions would be warranted while playing piano given that the first nine events corresponded to "notes" or "keys". Subject #4 said she thought of the three boxes as gauges on the panel of an airplane cockpit but could not explain how conceiving of them in that way led her to the correct answer. Subjects #2 and #5 could not articulate any well-formed statement about a strategy. They simply decided what the "right answer" was by "watching the flashes". They could not (or would not) say what it was about the flashes that provided information about the tenth event's location.

Note that all subjects reported using more than one strategy over the course of the experiment because no rule "worked" all of the time. Post-experimental interviews revealed that subjects felt frustrated in their attempts to "learn the rules" that determined the location of the tenth stimulus because rules or strategies that were successful at some point in the experiment were unsuccessful in others. All of the subjects reported that even during stretches of trials when rules were successful, they sometimes followed their "hunches" or predicted according to "what felt right". When rules started to fail, they fell back on these intuitive strategies in order to predict.

Data from the memory question in the first 500 trials showed serial position effects as shown in Table 4. Subjects did not exhibit any special sensitivity to the two crucial events in the memory task; the probability of a correct recall for events in those positions was not different from what would be expected in a "pure" serial recall task (Zechmeister & Nyberg, 1982). This supports an "unconscious" aspect of implicit learning in that subjects show no special memory for events that they must be processing given their observed levels of accuracy. Subjects #1 and #2 mentioned, however, that the presence of the memory question disrupted their "concentration" on the pattern of events, but, as Table 5 shows, average prediction accuracy did not change in the first 100 trials without the memory question (trials 501 through 601).

Discussion:

As noted above, the two rule changes in the last 1000 trials constituted a test of the subjects' knowledge of the pattern used in the first 2000 trials (Lewicki et al., 1987 and Nissen & Bullemer, 1987 used this technique to test their subjects' knowledge base). The findings above indicate that knowledge of the rule was at least robust enough to guide successful prediction behavior in Phase I in the absence of awareness of the rules or of learning.

Subjects' accuracy just after the rule shift was driven below chance due to response perseveration: subjects were continuing to predict according to a rule that had been very successful just the block before. Thus in the shifted rule condition, accuracy approached chance only after such prediction behavior ceased. Since Phase II consisted of only one-fourth as many trials as Phase I, the perseveration behavior kept subjects from reaching significant accuracy individually or as a group. Without the negative effect of perseveration at the beginning of Phase II, subjects might have learned the shifted rule and provided clearer evidence of knowledge transfer. Nevertheless, Phase I data indicated the implicit learning of an extremely remote, double-dependency rule.

During analysis of these results, however, an alternative analysis of the prediction data was discovered[1]. This alternative concerns the structure of the interdependencies among the crucial events and the feedback in the first 2000 trials. Specifically, knowledge of the location of either of the two crucial events alone can be of limited use in prediction because such knowledge precludes certain tenth locations and makes the remaining candidates unequally likely. For example: referring to Table 1, Column I, note that regardless of

---

[1] James McClelland proposed this alternative analysis in a personal communication in 1989.

the spot's seventh location, when the spot's third location is Box 1, its tenth location can be Box 1 with $p = 0.333$ or Box 2 with $p = 0.667$, but it can never be Box 3. Table 6, Column I, is a summary of prediction patterns that are implied by the use of such a "single location" prediction strategy involving either the third or the seventh spot's location. If a given subject used either rule, his or her pattern of predictions should conform closely to the prediction probabilities in Table 6, Column I. A subject using such a strategy exclusively could achieve a maximum accuracy of approximately 40% correct predictions as calculated by the following analysis.

Suppose a subject notices that if the third event is in Box 1, the answer is never Box 3, and if the third event is in Box 3, the answer is never Box 2. Suppose also that the subject uses three prediction rules: (1) Whenever the third event is in Box 1, predict Box 1 and Box 2 equally often; (2) When the third event is in Box 3, predict Box 1 and Box 3 equally often; and (3) Whenever the third event is in Box 2, guess. In this scenario, using rule (1) and (2) results in accuracy of 0.5(0.33) + 0.5(0.67) correct or 0.5 overall for each of the two rules. Using rule (3) results in accuracy of 0.33 overall. So for all cases when the third event is in box 1 or box 3, the subject will be correct on half the trials, and when the third event is in box 2 the subject will be correct on half the trials. The entire strategy thus predicts 0.165 + 0.165 + 0.109 = 0.439 correct. Using the "two location" prediction strategy, a subject could potentially achieve 100% accuracy.

Still other alternative rules exist. For example, Scarborough (1989) discovered the following strategy: first, consider either the third or the seventh event (the strategy works equally well for both), and then predict according to these rules: (1) If the chosen event is in Box 1, predict Box 3; (2) if it's in Box 3, predict Box 2; and (3) if it's in Box 2, guess. This strategy results in overall accuracy of 0.33(0.67) + 0.33(0.67) + 0.33(0.33) = 0.55 correct.

These alternatives raised the possibility that subjects were learning a less complex rule--one requiring procedural knowledge of one event instead of two, and perhaps even less complex than that. Examining prediction accuracy alone cannot reveal whether a one- or two-location strategy was learned; it is impossible to separate the contributions of the alternative prediction strategies to the data. Re-analysis of the data, however, indicated that although the subjects' overall accuracy level was indeed close to what the use of a single location strategy would predict, their actual patterns of predictions (which appear in Table 6, Column II) do not conform to the patterns in column I, suggesting that a single-event rule was not used, or at least not used exclusively. This is the only evidence remaining for the sensitivity of subjects to the double dependency. Additional research was deemed necessary to explore fully the possibility of recruiting implicit learning using this eclectic methodology, and it is described in the next chapter.

CHAPTER 3[2]

Experiment 2, presented in this chapter, was proposed with a view toward two principal theoretical goals. The first was to investigate further the depth and scope of implicit learning abilities, and so the methodology of Experiment 2 combined elements from previous experiments which were successful in recruiting implicit learning processes. The second goal was to link the probability learning studies, the pattern learning work of the Lewicki group, and the sequence learning work of the Cleeremans group more strongly to the grammar studies, the point being that all four areas need to be viewed as parts of a corpus of research lending convergent support to a unified explanation of implicit learning.

Clearly, these goals could not be attempted without a pragmatic effort to find a deterministic rule structure immune from the kinds of flaws embedded in the rule used in Experiment 1 and the rule used in Lewicki et al. (1987). Having accomplished that, the design of the experiment was kept similar to that of Experiment 1 in order to explore the rule-change and random-rule issues. In short, the flaw was "ironed out" in order to see

---

[2] The research in this chapter has been presented previously at the 62nd Annual Meeting of the Eastern Psychological Association, the 3rd Annual Convention of the American Psychological Society, and the 13th Annual Meeting of the Cognitive Science Society. The research was supported by National Science Foundation Grant BNS-89-07946 to Arthur S. Reber.

whether learning would be displayed if all other factors were held similar.

Experiment 2

Subjects:

Subjects were 6 Brooklyn College undergraduates. They were paid $40.00 for participating plus an incentive fee at a rate identical to that of Experiment 1 (one cent per correct response). They were run individually in small isolated rooms set aside for the purpose.

Apparatus and display:

The apparatus and display were similar to those used in Experiment 1 with five important changes. First, the three boxes were now located at the vertices of an invisible inverted triangle with Box 1 in the upper left-hand corner, Box 2 in the upper right-hand corner, and Box 3 in the bottom center of the monitor. Second, a trial was shortened to five successive events with subjects required to predict the sixth. Third, the duration of each appearance of the spot was shortened to 250 ms, the interstimulus interval was shortened to 250 ms, and the correct answer was displayed for 2000 ms. Fourth, the memory question was eliminated. Fifth, the sequences were produced using an exhaustive technique to ensure that no untoward sequence effects could contaminate the data.

Design and stimulus generation:

There are 243 (3^5) unique sequences of 5 stimuli distributed among the three locations, so that there are 243 possible patterns of 5 spots to appear successively in the 3 boxes. Eighteen different random orders of this exhaustive set of 243 patterns were constructed. Each of these 18 sets of 243 sequences was then blocked into 3 groups of 81 trials, for a total of 54 blocks. Subjects ran 9 such blocks a day for 6 days, for a total of 4374 trials. The nine blocks were really just 3 of the random orders broken into thirds (subjects were unaware of this); a block size of 81 was chosen based on pilot data which indicated that longer blocks often resulted in subject fatigue. Each subject observed the 18 random orders in the same sequence. There was a short pause between blocks.

The reasons for generating the stimuli in this manner rather than randomly, as they were in Experiment 1, are related to the nature of the new rule used in this experiment. The experiment consisted of three phases similar to those in Experiment 1. Phase I (the training phase) of the experiment consisted of 2430 trials, or 30 blocks of 81 trials. During this phase, the location of the 6th spot could be predicted with 100% accuracy according to the positions of the second and fourth spots in the immediately preceding sequence of 5 according to the following rule: If the second and fourth appearances were both displayed in the same box, the location of the sixth spot was Box 1; if they appeared in different boxes in a clockwise pattern (1 and 2, 2 and 3, 3 and 1 for the second and fourth spots, respectively) the location was Box 2; and if they appeared in different boxes in a counter-clockwise pattern (2 and 1, 1 and 3, 3 and 2), the location was Box 3. These rules are summarized in Table 7, Column I. Note that the confounds present in the rule used in Experiment 1 are precluded by the use of this "circular" rule system.

Phase II (the transfer phase) consisted of 972 trials (12 blocks of 81) in which the correct answer was, as in Experiment 1, shifted one box: patterns of the second and fourth spots which had predicted the sixth spot to be in Box 1 now predicted Box 2; those that predicted Box 2 now predicted Box 3; and those that predicted Box 3 now predicted Box 1 (see Table 7, Column II). Note that the number of trials in Phase II (972) is 40% of the number used in Phase I, a proportion greater than the 25% used in Experiment 1. The phase was lengthened in order to allow subjects more trials in which to learn the shifted rule.

Phase III, the random phase, consisted of 972 trials in which feedback to the prediction question was random, that is, no rule was in force to govern the displayed feedback (see Table 7, Column III).

Three principal reasons motivated the changes in which the stimuli were generated. The first was to eliminate coincidental "runs" of patterns that might allow subjects to construct elementary "local" rules independent of the underlying rule system that controlled the location of the target (6th) event. Second, this technique ensures that all subjects are exposed to the exhaustive set of patterns. This procedure guarantees that any observed accuracy cannot be accounted for by a preponderance of correct responses to one or two of the three kinds of trials, since each block of 243 contains an equal number of patterns that predict each answer. The third reason was to eliminate any potential bias introduced into a randomizing process via specific hardware or randomizing routines.

Once again, subjects were kept unaware of the existence of the three phases, and the following instructions were provided:

INSTRUCTIONS

On each trial of this experiment, you will see three boxes arranged as in the figure below [the inverted-triangle array was pictured below this text]. A spot will appear among the boxes 5 times. You should pay attention to where the spot appears because after the 5 spots appear, you will be asked to predict in which box (1, 2, or 3) the 6th spot will appear. After you predict, the correct answer will be shown on the screen and then go off; press the space bar to go on to the next trial. To make your prediction, use the 1, 2, and 3 keys at the lower right-hand corner of the keyboard.

You will be given practice trials before starting the actual experiment. If you wish, you may re-run the practice trials as well as re-read these instructions before any session to make certain you are familiar with the procedure. You will not be allowed to make notes while running the experiment.

You will run 729 trials a day each day for 6 days. Each days' trials are divided into blocks of 81 trials each. When you reach the end of a block of 81 trials, the computer will tell you "END OF THIS BLOCK". You can take a break at that point or go on to the next block. When you have completed all that day's trials, the computer screen will go blank. That is your signal that you are finished for the day.

If you have any questions about any of these instructions, please ask the experimenter.

No rule structure was in force during the 10 practice trials each subject observed before beginning the experiment proper, but all other attributes were identical to those in the rule-governed trials. Following Phase III, extensive post-experimental interviews were conducted with each subject. In addition to the open-ended interview, each subject completed the following ranking task:

Instructions: Rank the 5 appearances of the spot according to how important each one was in predicting the location of the 6th spot. Mark a "1" next to the spot appearance that was the MOST important in predicting the 6th spot's location, a "2" next to the appearance that was 2nd most important, and so on, down to a "5" next to the appearance that was the LEAST important.

Numbered spaces were provided for the subjects to mark their ranks. Finally, each subject was debriefed and paid.

Results:

Figure 9 shows the proportion of correct responses for each session of the experiment averaged across all subjects; individual data for each subject appear in

Figures 10 through 15. Once again, a line has been drawn for ease of comparing subjects' performance with the performance expected by chance alone. Block-by-block accuracy is reported in Table 8 for all subjects, and Table 9 presents the average accuracy, by subject, in each of the 3 phases.

As hypothesized based on the results of Experiment 1, subjects improve in their prediction accuracy over the training phase of the experiment, eventually reaching an average of 45.0% correct responses in the 10th block of 243 trials. Chance levels of prediction accuracy would yield 33.3% correct predictions. All accuracy figures over 38.7% are significantly above the level expected by chance alone in a block of 243 trials ($p < 0.05$). For Phase I as a whole (2430 trials), a given subject's accuracy must be at least 34.9% to be significantly above chance.

Similar to the results of Experiment 1, subjects' accuracy declines at the beginning of Phase II relative to the end of Phase I, but gradually increases. Average accuracy in the last 243 trials of Phase II is 42.1% correct predictions, which is significantly above chance ($p < 0.05$). Overall, subjects average 39.1% correct predictions in Phase II as a whole which is above the minimum 35.8% accuracy necessary for above-chance accuracy. The larger number of trials in this Phase relative to Phase II in Experiment 1 clearly contributed to subjects' significant prediction accuracy. Finally, as expected, prediction accuracy in Phase III drops relative to the last block of Phase II and never systematically changes.

During post-experimental interviews, all subjects reported initially trying to "crack the pattern" or discover some relationship among the pattern of 5 appearances, or some subset of the pattern of 5, and the feedback they were given. They reported intermittent success with such strategies: they would "work for a while" and then not work so well, but attempts to refine the strategies were not successful. The four subjects with the best accuracy (#2, #3, #4, and #5) reported that letting the pattern "wash over them" or letting the answer "come to them", instead of concentrating on where the five spots specifically appeared, seemed to produce the best results. Subjects #2, #4, and #5 all reported occasionally contradicting the answer suggested by such a strategy in favor of the more "mental" or reasoned choice; they noted that this usually resulted in an incorrect answer which would have been correct had they followed their intuition. As in Experiment 1, subjects could not specify any part of the rule which governed the feedback at any point in the experiment.

When asked specifically about the prediction strategies they used, subjects were unable to give well-formed statements of the rules in such strategies. They reported that it was difficult to remember all of the strategies that they had tried over the 6 days because they all seemed to "blend together", and what they remembered most was alternating their response strategies when one became too difficult to remember or "stopped working". All subjects mentioned that they were sure that "something had changed" about the rules at different points in the experiment, presumably alluding to the movements from phase to phase. But they were unable to say what it was that had changed or specify any particular attribute other than that the whole situation had been altered somehow. Constant probing seemed to make the subjects indignant about not having articulable rules, and typically, after 15 or 20 minutes of rigorous questioning, subjects eventually said "I don't know how I knew it" or "I can't remember now how I did it". Interestingly, findings from recent attempts to model this experiment with a parallel distributed processing (PDP) network shed light on subjects' inability to describe what elements of their knowledge were responsible for successful performance. These findings will be explored in more detail in Chapter 5.

Subjects' poor knowledge of the constraints embedded in the material was also confirmed by the ranking task, in which they were asked to rate each of the five stimuli in terms of their relevance in predicting the location of the sixth stimulus. The results failed to reveal sensitivity to the crucial events: on a scale of 1 (most important) to 5 (least important), the crucial events received average ranks of 3.5 (second event) and 2.67 (fourth event), whereas the first, third, and fifth events received average ranks of 3.33, 3.67, and 1.83, respectively.

In contrast, there was evidence that particularly salient sequences elicited high prediction accuracy even before subjects' overall accuracy reached significantly above-chance levels. For example, in Phase I, sequences in which all five stimuli appear in the same box (11111, 22222, and 33333) all predict Box 1 as the location of the sixth stimulus. Overall, subjects correctly predicted the successor of the three repeating sequences 61% of the time. They were also quite sensitive to patterns that alternated between two boxes (12121, 21212, 32323, 23232, 13131, and 31313), all of which also predict Box 1 as correct. They achieved 49% correct prediction accuracy overall for these sequences. These data indicate that, at least in some specific cases like these, subjects have become aware of some of the regularities embedded in the material. Table 10 shows each subject's probability of responding correctly to 27 patterns which were deemed to be highly salient due to their repeating, alternating, or circular movements. Note that subjects observe 10 exemplars of each of the salient patterns (one in each of the ten random orders in Phase I), so probability figures are only reported to one decimal point. This analysis revealed that only the first 9 patterns in Table 10, noted above, resulted in significantly above-chance accuracy, and these 9 were singled out for further analysis. The question of interest was what contribution the sensitivity to these patterns made to the subjects' overall observed accuracy. Subjects' successful prediction performance was not found to be based entirely on their sensitivity to these nine patterns: the average prediction score during Phase I dropped from 40.5% to 40.1% (still well above chance) when the 3 repeating and 6 single alternating sequences were eliminated from the analysis (Table 11). In order to discover <u>when</u> during the course of Phase I subjects became sensitive to these patterns, further analyses to test subjects' progressive sensitivity to the 9 highly salient were performed and appear in Table 12 and Table 13. Table 12 shows that there is no overall trend across subjects toward becoming sensitive to the salient patterns: <u>some</u> subjects learn <u>some</u> patterns "all at once" and subsequently predict correctly each time they appear.

Other subjects have only intermittent success with the 9 patterns and do not necessarily predict their successors correctly more in the latter stages than the early stages of Phase I. Table 13 presents the same data averaged across subjects over pairs of blocks. More of a trend toward increasing sensitivity is evident when the data are viewed this way, but the learning curve is still a bit jagged, especially for the alternating sequences.

Another observation concerns the data that subjects' prediction accuracy falls below chance levels when the rule which determines feedback is changed at the boundary of Phases I and II. A plausible hypothesis for this finding is that the below-chance accuracy is due to the subjects' continuing in their response patterns from the end of Phase I. This possibility was tested by comparing subjects' accuracy just before and just after the rule-change boundary between Phases I and II. Table 14 shows that if the Phase I rules were still in effect during the earliest part of Phase II, subjects' accuracies in the 81 trials just before and just after the rule-change are nearly identical. This indicates that response behavior at the very end of Phase I is continued in the early part of Phase II, contributing to the below-chance accuracy.

Further support for this last point is that the same kind of "carry-over" of response behavior is seen in the early stages of Phase III, immediately after the rule is changed for the second time. Table 15 compares subjects' actual prediction accuracy with what it would be if the Phase II rules were still in effect.

Discussion:

This experiment represents the best evidence to date that human subjects are sensitive to a remote, double-dependency rule, as measured by prediction accuracy. As noted in the introduction, there are at least two reasons why this task may be quite difficult from an objective, problem-solving point of view. The first is that there are three irrelevant events for every two crucial events, and the second is that each pair of crucial events is embedded in a large number of different, equally irrelevant contexts. This entails that subjects, inasmuch as they are successful at the prediction task, must extract the relationships among the crucial events and the feedback in spite of these two factors.

The results of Experiment 2 suggest that subjects are indeed succeeding in this task and that, moreover, the knowledge which guides their successful prediction behavior seems to be largely implicit: aside from the 3 repeating highly salient patterns, which all subjects seem to eventually notice explicitly, and the 6 alternating patterns, to which some subjects are sensitive, there is no conscious knowledge of the rule structure which governs the stimulus display. This is borne out in the results of the ranking task as well as the post-experimental interviews. Additionally, although there were no formal interviews after each day's sessions, some subjects remarked that, when the rule was shifted after Phase I, "something had changed" (or words to that effect) but that they did not know what it was. None used the word "rule" or "pattern". Similar comments were heard after the second rule change (between Phases II and III).

Data from the third (random) phase support the implicit knowledge point in a different way. The fact that prediction accuracy data from Phase III show no systematic, organized trend supports the idea that subjects are indeed using a rule in Phase II which is not "working" in Phase III. For the same reason, the data also serve as a control which suggests that there are no artifacts in the rule structure which allow subjects to exploit some unnoticed and unintended regularity.

Certain aspects of these results raised two issues--one old and one new. The old one was that, as in Experiment 1, during post-experimental interviews, subjects found it difficult to give an account of their strategies early in the experiment. Having run 4374 trials over six days, probing them about "what they knew and when they knew it" over the entire course of the experiment proved difficult, since the many identical trials seemed to "flow together" and were not clearly separable in their memories. And since they were not aware of the rule change boundaries until the debriefing, they could not accurately compare their performance in different phases. Interpretation of the ranking task results proved difficult for the same reasons. The task, like the interview, was conducted before the debriefing (for obvious reasons), so it was not clear from what part of the experiment subjects were basing their rankings. Subjects may have "lost" or forgotten knowledge that they possessed earlier about the relevant events. Another issue was raised with the observation that at the end of Phase I, subjects were still improving in their prediction accuracy--they had not yet levelled off at any clear asymptote. These two issues were addressed in Experiment III, along with the inclusion of an instructional variable.

CHAPTER 4

Experiment 3 was designed to address the issues raised at the end of the previous chapter by 1) increasing the number of trials to explore the issue of whether the learning was bounded at some upper level; 2) including daily tasks or interviews to assess each subjects' knowledge over the course of the experiment; and 3) exploring the effects of an instructional variable which sensitized some subjects to the existence of a rule. This last manipulation has precedent among other implicit learning studies, and they will be considered during the discussion of Experiment 3.

Experiment 3

Subjects:

Subjects were eight Brooklyn College students. They were paid $50.00 for participating plus an incentive fee of one cent for each correct response beyond the number expected by chance alone.

Apparatus and display:

Apparatus and display were identical to that of Experiment 2.

Design and stimulus generation:

The entire experiment consisted of 12,150 trials. Each trial was identical to the trials in Experiment 1, but there were no distinct phases and no rule shifts of any kind. Thirty-two additional random orders of the set of 243 possible patterns were constructed for a total of 50 orders, or 12,150 trials. Each order was again blocked into thirds to achieve blocks of 81 trials each; each subject ran 9 blocks (729 trials) a day for 16 days; on the 17th day they ran 6 blocks (486 trials). Finally, response latencies (in milliseconds) were collected on each trial beginning with the appearance of the prediction question on the monitor and ending with the subject's response. Subjects were not informed that response latencies were being collected.

Instructions:

The subjects were randomly assigned to one of two conditions called "neutral" and "explicit". Neutral subjects received instructions identical to Experiment 2 subjects. Explicit subjects received the same instructions but with the following paragraph inserted between the first and second paragraphs of the Experiment 1 instructions:

> Before you begin, you should know that there is a rule that dictates where the 6th spot will appear. As you know, you are being paid an incentive rate for each correct prediction you make. Therefore, if you can discover the rule that dictates where the sixth spot will appear, your incentive payment will be higher than if you cannot discover the rule.

Both versions of the instructions were altered to reflect the larger number of trials and the running schedule noted above.

Daily tasks:

After each days session, neutrally instructed subjects were asked to summarize their strategies for predicting during that day's session, indicating any changes from the previous day's session. They were also encouraged to state anything new or unusual that they had noticed or discovered. The interviews were informal and unstructured, allowing the subjects to state, with whatever vocabulary seemed most comfortable, exactly how they chose their answers to the prediction questions. No feedback of any kind was given to the subjects' remarks. Explicitly- instructed subjects were asked to complete the ranking task from Experiment 1 after each day's session. They were not asked for the same type of daily summary as the neutrally-instructed subjects since it was assumed that their efforts were basically aimed at discovering the rule which they knew existed.

After the last session, each subject in both conditions performed a prediction task consisting of 30 sequences identical to ones in the experiment except that no feedback was given; subjects had to state their prediction verbally and justify it as concisely as possible. The purpose here was to discover what strategies or rules subjects were using during the experiment. This task seemed a more fundamental way of doing that rather than simply asking them to report directly the rules they used, although this was also done after the predict-and-justify task.

The 30 sequences were chosen to reflect a cross-section of the exhaustive set of patterns that also contained the patterns that were known (from Experiment 2) to be highly salient. Thus they were selected by picking one of the 50 random orders at random, starting at the top of the list of 243, and choosing the first 30 patterns with the stipulation that they be divided equally among the three types of patterns (same box, clockwise, and counter-clockwise). Once 10 of each pattern had been chosen, similar exemplars of that category were rejected except for the sub-requirement that only one of each of the following four patterns was included: continuous (e.g., 22222), single alternating (e.g., 32323), purely clockwise (e.g., 12312), and purely counter-clockwise (e.g., 32132) (Table 16 is a complete list of the 30 patterns in the order in which they were presented). The first stipulation was included to ensure that, as in the experiment itself, subjects observed an equal number of each of the three types of patterns. The second requirement was included so that subjects' sensitivity to those four specific pattern types could be assessed while not including more than one of each, which might have led some subjects to construct local rules. Finally, each subject was debriefed and paid.

Results:

In general, the results from the neutrally-instructed group replicate the finding from Experiment 2 that subjects learn the dependency embedded in the display but cannot articulate explicit knowledge sufficient to justify their prediction accuracy (as noted above, the behavior of PDP models trained in a similar fashion is relevant here and will be explored in the next chapter). But a broad overview of either of the groups is misleading because unlike Experiments 1 and 2, some subjects in Experiment 3 actually discovered the rule, and others showed no evidence of learning whatsoever. Because deeper analyses are necessary in order to discover "what's going on" in this study, this section is divided into several sub-sections in order to deal with the large amount of data presented. In some cases, the neutrally- and explicitly-instructed groups are discussed separately within each sub-section. In all figures and tables, the symbol 'E' or 'N' next to a subject's number refers to the explicit or neutral instructions, respectively, that he or she received.

I. Overall accuracy.

Explicitly-instructed subjects:

Average prediction accuracy for the explicitly-instructed subjects appears in Figure 16 (bottom), and reflects the assumption that the two subjects who discovered the rule (3(E) and 5(E)) would have achieved 100% accuracy for each session after the discovery. Individual accuracy data are in Figures 17 through 20. The figures for the two subjects who discovered the rule (3(E) and 5(E)) reflect only the sessions up to and including the session in which they discovered the rule. Accuracy is reported as a function of the 17 daily sessions (note that all sessions except the 17th represent 729 trials; the 17th session represents 486 trials). Once again, a horizontal line appears in each figure at the level of accuracy expected by chance alone (33.3&). Block-by-block accuracy data for each subject appear in Table 17, and average block-by-block accuracy appear in Table 18. Table 19 displays average prediction accuracy for each subject.

Still considering the explicitly-instructed subjects, it is instructive to examine the attributes of the subjects who did and did not discover the rule which governed the stimulus display.

Rule-discoverers:

Subject 3(E) discovered the rule on the 9th day of running, subject 5(E) on the 5th day. In the three random orders of the exhaustive set of 243 patterns they ran on those days, subject 3(E) correctly predicted the sixth event 131 times, 237 times, and 242 times; subject 5(E) correctly predicted 229 times, 234 times, and 232 times.

Before being told these results or indeed anything about the experiment or their other results, they were asked to write the rule down, and here are their transcripts:

Subject 3(E):
The base rule is to ignore the first, third, and last flash. Use the second and forth [sic] in the following manner: if both the second and forth flash is [sic] in the same box than [sic] the answer is 1. If the forth is mathematically after the second then the answer is 2.

Ex.: 2nd = 2 4th = 3 Answer = 2
2nd = 3 4th = 1 Answer = 2
2nd = 1 4th = 2 Answer = 2

If the forth is mathematically before the second then the answer is 3.

Ex.: 2nd = 3 4th = 2 Answer = 3
2nd = 2 4th = 1 Answer = 3
2nd = 1 4th = 3 Answer = 3

Note: Rather than go higher that 3 or lower than 1 we simply jump to the end or beginning of the loop.

Subject 5(E):
The 2nd number determines what the 6th number will be by virtue of the 4th number. If the 2nd number is a 1 then the 6th number will be whatever the 4th # is. If the 2nd # is a 2 then the 6th # will be whatever the 4th # is - 1. In the case of the 4th being a 1 the 1 turns into a 3. In the case of the 2nd # being a 3 the 6th # will be 1 more than the forth [sic]. If the 4th # is a 3 it turns into a 1.

Both sets of rules accurately reflect the relationships among the two crucial events and the feedback, though obviously not in the geometric way conceived of in the design of the rule. Note that both rule-discoverers were explicitly-instructed and that neither of them discovered the rule within the first 2430 trials--which constituted the entire exposure to the rule for Experiment 2 subjects.

Rule non-discoverers:

As in Experiment 2, a block size of 243 trials requires a subject to achieve at least 38.7% accuracy in order to be significantly above chance; in the entire experiment (12,150 trials), accuracy must reach 34.0% for a given subject. Of the two explicitly-instructed subjects who did not discover the rule, one (1(E)) never reached significantly above-chance prediction accuracy throughout the experiment; her overall prediction accuracy in the entire experiment was 33.1%. The other (7(E)) gradually improved in accuracy, much like the subjects in Experiment 2, and ended up averaging 42.2% accuracy overall. This wide variation in performance among the explicitly-instructed group--from discovering the rule to showing no learning at all--will be discussed in more depth later.

Neutrally-instructed subjects:

Average prediction accuracy for the explicitly-instructed subjects appears in Figure 16 (top); Individual accuracy data are in Figures 21 through 24. Block-by-block accuracy appears in Table 18 (average) and Table 19 (individual). Among the neutrally instructed subjects, one (subject 6(N)) never reached above-chance levels of accuracy, achieving only 32.8% accuracy overall. The remaining three (subjects 2(N), 4(N), and 8(N)) improved gradually, reaching 43.4%, 39.5%, and 39.6% accuracy overall, respectively.

II. Accuracy in predicting highly salient patterns.

As in Experiment 2, subjects' accuracy in predicting the 6th event in the 9 highly salient patterns was assessed, and these data are presented in several ways. Note that since the analyses were aimed at assessing the relative contribution of the accuracy on these 9 patterns to the subjects' overall accuracy, data from the two explicitly-instructed subjects who discovered the rule were not analyzed. Table 20 displays each subject's prediction accuracy to the 9 highly salient patterns, and Table 21 displays the same data averaged across neutrally-instructed subjects, non-solving explicitly-instructed subjects, and all subjects. Table 22 presents the number of times each response (1, 2, or 3) was given to each pattern. Table 23 presents the progressive sensitivity of each subject to each of the 9 patterns by showing the probability of a correct response in groups of 5 consecutive blocks (exhaustive sets) of trials. Note that each subject observed 50 exemplars of each pattern, one in each of the 50 random orders.

Once again, subjects show high sensitivity overall to the 9 patterns, averaging 74.0% correct predictions to the 3 continuous patterns and 51.0% correct predictions to the 6 single-alternation patterns. In the 3 continuous patterns, the 4 neutrally-instructed subjects achieved 80.0% correct predictions and the 2 non-solving explicitly-instructed subjects averaged 79.6%--practically identical accuracy. In the 6 single-alternating patterns, however, the figures were 54.0% for the neutrals and 45.8% for the explicits. Not surprisingly, the two subjects whose accuracy never reached significantly above-chance levels (1(E) and 6(N)) show poor prediction accuracy relative to the other subjects for these patterns.

The progressive accuracy data in Table 23 show that the 4 subjects who did gradually improve over the course of the experiment learned the 3 continuous patterns within the first 15 blocks and never erred in predicting the successor to those patterns after that point. Accuracy in predicting the successor to the 6 single-alternating patterns do not show an analogous dramatic trend.

Overall prediction accuracy was compared with it would be if all exemplars of these 9 patterns were removed from the analysis of accuracy. Table 24 displays this analysis for the six subjects who did not discover the rule. Prediction accuracy of the 4 subjects whose accuracy improved over the course of the experiment fell from 38.4% overall to 37.6% overall, which is still above the level necessary for significantly above-chance prediction.

III. Response latency.

Response latency (in milliseconds) was measured in each trial from the time the prediction question appeared on the screen until the subjects made a response. Latency data for each of the 6 subjects who completed the experiment were analyzed. Table 25 presents the means and standard deviations of the data divided by accuracy and correct response.

Individual latency data were analyzed with t-tests on the means of correct and incorrect responses. The 4 non-solving subjects whose accuracy increased over the course of the experiment (subjects 2(N), 4(N), 7(E), and 8(N)) responded faster when they were correct than when they were incorrect: obtained $\underline{t}$s(12,148) were 12.05, 6.67, 5.91, and 4.96, respectively. The differences between the means are significant ($\underline{p} < 0.01$, two-tailed) for those 4 subjects (means appear in Table 25).

IV. Assessment tasks.

A. Ranking task (explicitly-instructed subjects).

Results of the daily ranking task for the explicitly-instructed subjects appear in Table 26 as the average rank assigned to each event after each daily session. These rankings reveal two similarities of the subjects who discovered the rule (3(E) and 5(E)): first, they became sensitive early on to the relevance of the second event and the irrelevance of the fifth event. Second, even in the last session before discovering the rule, they nevertheless both ranked the crucial event locations 5 and 2, respectively.

The subject who gradually learned (7(E)), however, did not display increasing sensitivity to the two crucial locations, ranking them 1, 2, 3, 4, and 5, respectively, in the last 6 sessions. Subjects 1(E) (a non-learner), less surprisingly, likewise showed no progressive sensitivity to the crucial locations. The bottom part of Table 26 shows that when the average overall rankings are considered, subjects do seem to know that events one and five are <u>poor</u> predictors, but they cannot discriminate among the middle three events' usefulness.

B. Daily interviews (neutrally-instructed subjects).

Daily interviews were conducted with each neutrally-instructed subject after each day's session. All subjects except 6(N), the non-learner, reported trying to "recognize patterns" in order to predict correctly. These efforts were only partly successful; as subject 4(N) put it, "you can only go so far with a rule until it doesn't work anymore". By that, she said that she meant that "some rules would work sometimes and not work other times". Subject 2(N) mentioned that in trying to refine a rule, she would "get so far away from where the rule started that [she] lost track of how to use it." By the third session, two subjects (2(N) and 4(N)) mentioned the continuous patterns and that they tried to induce from them a general a rule for predicting the 6th event. Both subjects could clearly report that "when all 5 spots are in the same box, the right answer is Box 1". They were pressed to state an analogous rule for some other pattern or set of patterns but could not.

Recall that data from Subject 6(N) give no evidence of learning. In the daily interviews for every session after the first, he indicated that he used one strategy exclusively to predict the target location. He reported that the feedback was unrelated to the patterns of five spots but instead appeared in runs, so that if a run where, say, Box 1 feedback was in effect, responding "1" during that period would be sufficient to achieve above-average correct response rates regardless of the pattern of five spots in that trial. He explained that while a run of a certain feedback was in effect, he would note where the feedback occurred in those trials where he was incorrect, because that determined where the next run would appear when the current run ended. Examining his data bear this out: there are long runs of from 10 to 40 trials during which he restricted his response to one key. It is puzzling to note that he persisted in this strategy when all of the feedback should have convinced him it was not effective. When asked about this post-experimentally, he maintained that this strategy <u>had</u> been effective and was surprised at his accuracy data. Apparently, he had succumbed to relying on small samples of trials in which he had been right more than half the time. His inability to be able to consider the long-range net effectiveness of his strategy led him to think he was performing at an above-average rate and thus entailed its continued use. Additionally, his accuracy data fail to show the high sensitivity to the 9 salient patterns compared to those who did improve over the course of the experiment, and his latency data indicate faster correct and incorrect responses than any other subject in either group. Unlike other subjects (in both groups) whose accuracy improved over the course of the experiment, the difference in his mean latency to respond both correctly and incorrectly is not significant.

Subject 8(N) reported using a "counting rule" which involved counting the number of events in each box and using that information to predict where the 6th event would occur. The rule was expressed as "the number of flashes in each box dictates where the 6th flash will occur." The subject used the single-alternating patterns to exemplify how the strategy worked, saying that when the flashes "flip-flopped" between two boxes, the answer would follow the alternating pattern, so that the pattern 12121 would predict 2 (which is incorrect), and 21212 would predict 1 (which is correct). However, his progressive sensitivity to those patterns does not indicate that the rule helped his accuracy to those patterns. Like subjects 2(N) and 4(N), he could not, when pressed, articulate a clear statement of the rule that produced testable strategies.

C. Predict-and-justify task (all subjects).

Accuracy results from the predict-and-justify task appear in Table 27. Subjects were encouraged to justify their responses as fully and with as much detail as possible. All of the subjects seemed to overgeneralize their knowledge of single-alternating patterns. When the spot appeared <u>mainly</u> in two boxes, but not necessarily in a regular alternating pattern, subjects usually justified the prediction by saying that the spot alternated between two boxes. To patterns that did not have a clear alternation, subjects often said "it looks familiar" or "I recognize that pattern". All subjects could not justify their responses to at least 5 or 6 patterns; they simply shrugged and reported that they "knew it somehow". When pressed on these occasions, they reported that it was "the way the spot moved" or the "character" of a particular pattern that guided their choice, but they could not give an algorithmic account of their decision process. Subject 7(E) reported concentrating on trying to see each pattern's "gestalt" properties, even to the point of defocussing her eyes in order to perceive it optimally.

She reported that at times throughout the experiment she would try to discover a rule by testing hypotheses but that she ultimately always fell back on the more intuitive strategy because it resulted in higher accuracy. She gave corresponding justifications to her predictions, mentioning that it was the way the spot moved or the "flow" of its particular path that dictated where its 6th appearance would occur.

A Pearson *r* was calculated for the 6 subjects who completed the experiment, comparing their overall accuracy in the experiment with their accuracy at the predict-and-justify task. The correlation coefficient is 0.55.

Discussion:

The general finding from Experiment 2, that subjects can learn a contexted, double-dependency rule, was replicated in half of the subjects: one explicitly instructed subject (7(E)) and three neutrally instructed subjects (2(N), 4(N), and 8(N)) seemed to acquire implicit knowledge of the rule structure, much like the Experiment 2 subjects during Phase I.

Data from several studies of implicit learning (Brooks, 1978; Howard & Ballas, 1980; Reber, Kassin, Lewis, & Cantor, 1980) show that explicitly instructed subjects in implicit learning experiments show poor performance relative to their neutrally-instructed cohorts. This is because the rules underlying the stimulus displays are so complex that overt, hypothesis-testing attempts to learn them are largely unsuccessful. As Reber (1989) phrased the issue, "looking for rules won't help if you can't find them". In this experiment, however, two explicitly instructed subjects did discover the rule governing the display, and it is interesting to note that they characterized it as an arithmetic algorithm. When these two subjects were debriefed, they were surprised that the rule (as originally construed) was spatial, and not mathematical, in nature. That fact does not diminish the applicability of the rule they articulated of course, but it does afford some insight into how they approached the experiment, given that they knew a rule existed--as opposed to the neutrally-instructed subjects who, while they may have suspected, did not know that a rule was in force.

Additionally, their block-by-block prediction data (Table 15) do not show progressive increases in accuracy leading to perfect prediction behavior. Until the session in which they discovered the rule, their prediction accuracy is comparable with other subjects'. When asked to account for his solution, subject 3(E) said that it was one of various prediction strategies he was testing early in the fifth session, that it seemed to work pretty well, and that he gradually refined the arithmetic combination until he discovered the rule. Subject 5(E) reported that he "invented" his rule while concentrating intently on trying to discover the relationship between the events and the feedback. Both subjects reported that once they realized that only some of the events were important, it was only a matter of isolating the nature of their relationship before the rule was discovered. The relationship could be figured out relatively quickly, they said, because there were enough trials in one session in which to gradually refine the rules they ultimately reported.

In summary, the neutrally instructed subjects in this experiment gradually acquired knowledge of the rule structure underlying the pattern without reporting knowledge of the rules or awareness of having learned anything. In the daily post-experimental interviews, subjects typically commented that the way in which the stimuli were presented seemed to induce them to concentrate less on the specific spot locations and more so on the set of five appearances as a whole. Related to that issue is that like subjects in Experiments 1 and 2, some subjects reported that they felt that responding according to their intuition was often more effective than attempting to refine rule-based prediction strategies which often became too cumbersome to be used profitably.

Related to this last point is an interesting hypothesis about the 3 continuous patterns, which subjects learn relatively early on and spontaneously mention in post-experimental interviews: early discovery of these patterns might actually work against learning the rule in that it might encourage a subject to use all five events as a basis for constructing potential rules. Following up this line of logic, the other 6 patterns could also be potential impediments to rule discovery if subjects generalize incorrectly about some aspect of the single-alternating patterns. This idea, not specifically addressed in this dissertation, could be explored with a future experiment which required subjects to give verbal protocols as they performed the task, thus providing insight into how their prediction strategies were based on the sequences to which they were exposed.

The explicitly-instructed subjects show larger individual differences than the neutrally instructed subjects. Besides the two rule-discoverers, one subject learned progressively and one seems not to have learned at all. The latter subject (1(E)) was, as noted above, generally faster to respond. Perhaps trying to rush through completing each daily session prevented her from concentrating on the display as much as was necessary to learn the embedded interdependencies. Unlike the other non-learning subject (6(N)) she did not report using an incorrect rule, and so there is no obvious explanation for her failure to learn. She did show outward signs of "rushing through" each daily session, however, and seemed to want to be finished with the experiment as quickly as possible. Although not rigorously scientific, a plausible explanation may lie in her simply not attending to the stimulus display long enough nor with sufficient concentration to learn.

One other observation stands out as regards the overall performance of the 4 subjects who learned but did not discover the rule (2(N), 4(N), 7(E), and 8(N)). Recall that at the end of Phase I (the learning phase) of Experiment 2, subjects had attained an average prediction accuracy of approximately 44%. But after an equal number of trials in Experiment 3, these 4 subjects had only attained 37% average accuracy; it is not until some point between 8500 and 9000 trials, depending on the subject, that they reach the 44% accuracy mark. A possible explanation for this slower learning is that the Experiment 3 subjects, knowing the length of the experiment, were not as motivated to learn at as high a rate as their Experiment 2 counterparts, since Experiment 3 is approximately three times the length of Experiment 2.

The next (and final) chapter is an overall discussion of Experiments 1, 2, and 3, the connections among them, and their relationship to cognitive psychological phenomena in general.

CHAPTER 5

This final chapter is a general discussion of the new research presented in this dissertation, its relationship to research discussed in Chapter 1, and its relevance to current issues within the field of implicit learning in general. Recall that one of the stated theoretical aims for pursuing the research presented here was to explore implicit learning theory with a novel methodology; the other was to keep the research comparable to the work of Reber, Cleeremans, Lewicki, Mathews, and their colleagues, cited here extensively, which forms the current theoretical foundation for that theory. To the extent that the second aim has been attempted, the current work must address the controversies and topics now concerning the field, explored in detail in a review article by Reber (1989).

A logical starting point for this discussion is a consideration of, first, the knowledge base subjects possess as a consequence of their participation in these experiments and, second, the process by which that knowledge "got there". Put simply, it is important to consider <u>what</u> subjects learn in this task, and <u>how</u> they learn it. The explanations are closely related. In the present experiment, significant levels of prediction accuracy are taken as prima facie evidence that learning occurred. But subjects, in all of the various probe techniques used here, find it difficult to articulate any explicit knowledge about the rules of the stimulus environment that they are clearly exploiting given their accuracy. For example, they show no special sensitivity to the crucial event locations on a ranking task. When required to justify their predictions to a representative cross-section of patterns, they do not articulate clear rules or strategies that generate testable prediction strategies. And when directly asked to relate rules, or parts of rules, that governed the stimulus display, they cannot articulate any coherent, well-formed statement concerning rules about how the target event was predictable. For all of these reasons, the subjects' knowledge seems to be largely implicit, that is, unavailable to conscious inspection.

This last point is only a partial response to the question of <u>what</u> subjects learn in these experiments. It begs the question of, assuming that subjects' knowledge <u>is</u> implicit, the structure and form of that implicit knowledge. Note that this is a different question: it addresses not the implicit or explicit nature of the knowledge but how that knowledge is represented in this implicit store. A comprehensive answer to the question of how even conscious, explicit knowledge is represented in memory has not been formulated by cognitive psychology, and the situation within the field of implicit learning is similarly far from clear. The traditional answer to this question is that the knowledge is abstract in nature: it is represented not as specific, concrete rules or exemplars, but as an abstract representation of the stimulus environment to which the subjects were exposed. Once again, this is not to say that all knowledge used by the subjects is abstract or even unconscious: the finding that many subjects notice highly salient patterns, can predict their successors with considerable accuracy, and mention them specifically when asked, is a testament to the fact that some part of their knowledge is indeed conscious and instantiated. Results like these have prompted Brooks (1978) to argue for an instantiated, non-abstract memory system that he claims can account for the results of artificial grammar studies. This dissertation has endeavored to show that the present subjects' performance cannot be fully accounted for on the basis of such instantiated knowledge alone. It seems reasonable to conclude that the amount of implicit knowledge they possess about the stimulus environment is greater than the amount of explicit knowledge they possess.

This conclusion, however, is also not without its critics. Dulany, Carlson, and Dewey (1984, 1985) dispute it as it relates to artificial grammar studies. They argue that bits of explicit knowledge about the underlying stimulus display (such as certain pairs of permissible bigrams), while incomplete in themselves, are enough to account for many of the learning effects reported in such experiments. Perruchet et al. (1990) used a similar analysis to rebut the conclusions of Lewicki et al. (1987). The rule used in Experiments 2 and 3 presented here precludes obtaining significant prediction results from the application of such rule fragments or bits of knowledge. Recall also that when subjects' fragmentary, explicit knowledge of highly salient patterns is discounted from analysis, their resulting accuracy is only barely changed.

What is the basis for the claim then, that at least in these experiments, the structure of the implicit knowledge is abstract in nature? Many of the early probability experiments of the sort discussed in Chapter 1 (e.g., Reber, 1967; Reber & Millward, 1968) produced considerable evidence that subjects in those studies induced a mental representation that was reflective of the structure of the stimulus environment to which they were exposed, but not of its manifest physical attributes. In these experiments, subjects typically observed long runs of simple events that followed probabilistic schedules of appearance and were then required to predict subsequent events. Millward & Reber (1971) found that subjects predicted accurately to the point of anticipating the changing probabilities of events, even when the anticipatory response required sensitivity to information integrated over the 50 preceding events. These subjects, rather than shadowing the fluctuating event likelihoods, learned to anticipate the shifts in event probabilities so that their predictions mirrored the sawtooth pattern of event probabilities over the course of the experiment. In another instance (Millward & Reber, 1972) subjects exploited a "flaw" in a computer program that displayed two events with probabilities of 0.520 and 0.480, instead of 0.500 and 0.500 as was intended. The subjects predicted almost perfectly (0.523 and .477) according to the odd, but consistent, event appearance-rates they observed in the resulting stimulus display (Millward & Reber, 1972). In these cases (and others; see Reber, 1989 for a review), subjects learn the underlying, probabilistic structure of the stimulus environment and are capable of exploiting it to meet task requirements. The case for an abstract representation of knowledge is clearest in these experiments; there is no convenient way to account for subjects' performance in terms of exemplars or an instantiated, explicit memorial system of any kind. Additional evidence for the abstract nature of subjects' knowledge comes from artificial grammar studies (Manza & Reber, 1991; Mathews et al., 1989; Reber, 1969). Subjects who are "trained up" on exemplars generated with one set of letters, but tested for their knowledge of the underlying rule structure with exemplars generated by the same grammar with different letters, can "transfer" their knowledge about the common underlying grammar and accurately discriminate grammatical from nongrammatical strings.

There is further support for the abstract nature of implicit knowledge from a very different source. Kushner, Cleeremans, and Reber (1991) "trained" parallel distributed processing (PDP) networks under exactly the same conditions as the human subjects in Experiment 2. The networks showed steadily increasing prediction accuracy over the course of Phase I, eventually reaching a level of accuracy almost identical to the humans'. The interesting finding regarding the abstractness issue was that the representations developed by the network were completely opaque--the information encoded by the network allowed successful performance, but was not readily decomposable into critical features. In fact, highly complex analyses, developed only very recently, are necessary in order to detect the regularities embedded in the internal representations developed by the network (see Cleeremans, Servan-Schreiber, & McClelland, 1990 for details of the analyses). This characteristic of the representational systems of PDP models in general provides a natural explanation for the fact that human subjects are unable to describe what elements of their knowledge are responsible for successful performance. Whether PDP models are to be taken seriously as literal reflections of human representational systems is arguable, but this provocative interpretation of the model's behavior speaks directly to the abstract vs. instantiated issue.

The conclusion about the abstract nature of implicit knowledge is not qualitatively different when experiments with more complex environments are considered. As mentioned in the Introduction, this process has been seen in widely different experimental situations that nonetheless share the property of using complicated stimulus displays which embed remote regularities. In the Lewicki et al. (1987) study, the rules, events, physical aspects, and even task requirements of the experimental situation were more intricate than those used in the probability learning experiments just considered. But the same arbitrary relationship among events was preserved, and with extended practice, subjects showed analogous improvement in performance. The conclusion here is the same as it was for the probability learning experiments: subjects appear to have induced a representation of the remote event interdependencies built into the stimulus environment which guides their behavior. Not surprisingly, they also show the expected lack of explicit, articulable knowledge about that environment. Experiments 2 and 3 presented here are comparable in that subjects seem to be extracting information about the deterministic rule structure which most of them cannot articulate but which guides their prediction behavior. Put another way, they possess procedural knowledge that seems to not be consciously available. Following this reasoning, Cleeremans and McClelland's (1990) failure to find evidence that subjects can be sensitive to crucial events separated by more than three elements might be due to their use of a probabilistic rather than a deterministic rule structure, although only a controlled experiment could fully develop this idea.

Reber (1989) cites studies which show that explicit attempts to learn in complex, rule-governed situations like the present one degrade performance compared to the use of strategies that invoke implicit processing. Two subjects in Experiment 3, however, did discover the rule, and interestingly, they both received explicit instructions. Note that since the total sample of solvers is only 2, the likelihood of getting both rule-discoverers in the explicit group is 0.25. Keeping in mind the caveat of not drawing too-strong conclusions with studies with small Ns, it seems that the structure-extraction process proceeds best (for most individuals) implicitly, but that individual differences may favor explicit hypothesis-testing for some individuals. There is some precedent within the field of implicit learning for individual differences when subjects are encouraged to explicitly search for rules in experiments designed to invoke implicit learning (Reber, 1976), and it seems to be repeated here.

Some very recent work may explain the success of the two subjects who discovered the rule--two instances of an event <u>without</u> precedent in any reported implicit learning experiment. Reber, Walkenfeld, & Hernstadt (1991) found IQ, as measured by the WAIS-R, to be more highly correlated with performance in explicit tasks than implicit tasks. Experiment 3 contained no a priori hypotheses about this issue, and IQ was not measured. But personal communications from the two subjects' teachers reported them to be extremely bright, motivated, and generally intellectually gifted. This admittedly anecdotal evidence may hold the explanation for this surprising result, though additional research would be necessary to establish the explanation as sound. In general, more detailed studies devoted exclusively to investigating explicit learning abilities are necessary in order to fully explore this issue.

An additional question within the knowledge representation issue concerns the psychological representation of the rules built into the stimulus environments. The question stems from the fact that there is a potentially infinite number of rule structures, many (or even most) of which have no resemblance to that which the experimenter conceived, that could generate identical "correct" answers to each of the 243 patterns used in these experiments. The arithmetic characterization by two subjects of what was conceived as a spatial rule in Experiment 3 is a good case in point. Reber (1989) uses the analogy of transformational grammar as an example of a very different characterization of grammar than the one most people learn or know to be accurate, but which is nevertheless a legitimate formalization of syntax. Assuming that subjects have implicitly extracted an abstract representation of the stimulus environment, we cannot assess the nature of that representation since they are unable to articulate any part of its structure. To revisit a venerated proposition, behavior, in this case prediction accuracy, is the only unassailable evidence for their knowledge--knowledge, perhaps, of the rule imposed on the stimuli by the experimenter, but knowledge of some rule, or set of rules, in any case.

The criticisms of Lewicki et al. (1987) by Perruchet et al. (1990) speak to this issue, as does the flaw in the rule used in Experiment 1 reported here. In both cases it was shown that a "simpler" rule, subset of rules, or strategies unrelated to the built-in rules, might account for the observed performance of subjects. Perhaps, arguing from the points in the last paragraph, even simpler strategies are still lurking, undiscovered, among the possibly infinite alternative characterizations of the rule structures. But the force of these findings does not diminish the point that in both cases, subjects did extract, implicitly, procedural knowledge of the respective stimulus display. Regardless of how many simpler rules or strategies exist, the finding that learning occurred implicitly is not affected. Similarly, the knowledge base resulting from this learning, irrespective of its contents, is still not available to conscious inspection.

Following up on this point, it is interesting to note that Perruchet et al. (1990) acknowledge the possibility of both unconsciously held knowledge and of the process of abstraction of a rule structure, yet they argue forcefully against the possibility that abstraction can occur without consciousness of the process. Their position seems to be that when unconscious processes can be shown to be operating, there is no indisputable evidence for abstraction, and conversely, that when there is sound evidence for abstraction, consciousness will be seen to be operating. The best rejoinder to this position seems to be that no claims for the exclusive operation of implicit processes have been made for the present set of experiments. Transcripts of subjects' post-experimental comments alone reveal that both implicit and explicit processes are involved when they confront this task, although in a complex, interwoven way. Although the majority of neutrally-instructed subjects report eventually abandoning "purely" explicit learning processes, it does not seem controversial to claim that such processes are never completely absent. Human beings are inveterate problem solvers, constantly searching for patterns in the environment even when those patterns are deeply remote or even absent (Reber, 1989; Simon & Sumner, 1968). The crucial point to be concluded from these experiments is that where patterns too remote to be detected explicitly do exist, implicit processes seem to be automatically invoked and utilized in the service of learning. They produce a virtual representation of the structure of the rule environment which is unavailable to conscious inspection but which, along with consciously held knowledge, guides behavior.

LOCATION OF SPOT ON THIRD APPEARANCE:	LOCATION OF SPOT ON SEVENTH APPEARANCE:	LOCATION OF SPOT ON TENTH APPEARANCE: I DURING TRAINING ([illegible] TRIALS):	II DURING SHIFTED FEEDBACK CONDITION (900 TRIALS):	III DURING RANDOM FEEDBACK CONDITION ([illegible] TRIALS):
BOX 1 -------->	BOX 1			
BOX 2 -------->	BOX 2	BOX 1	BOX 2	BOX 1, 2, OR 3 p = 0.33
BOX 3 -------->	BOX 3			
BOX 1 -------->	BOX 2			
BOX 1 -------->	BOX 3	BOX 2	BOX 3	BOX 1, 2, OR 3 p = 0.33
BOX 2 -------->	BOX 3			
BOX 3 -------->	BOX 2			
BOX 3 -------->	BOX 1	BOX 3	BOX 1	BOX 1, 2, OR 3 p = 0.33
BOX 2 -------->	BOX 1			

Table 1: Summary of rules governing feedback in Experiment 1.

Subject:

	Block	1	2	3	4	5	Mean
	1.	.390	.325	.340	.335	.345	.347
P	2.	.355	.335	.345	.320	.345	.340
H	3.	.405	.365	.385	.325	.325	.361
A	4.	.440	.340	.325	.385	.340	.366
S	5.	.400	.345	.365	.365	.360	.367
E	6.	.410	.365	.355	.375	.365	.374
	7.	.395	.370	.370	.370	.405	.382
1	8.	.425	.390	.375	.385	.395	.394
	9.	.390	.395	.390	.405	.415	.399
	10.	.390	.391	.400	.395	.405	.396
		1 Block = 200 Trials in Phase 1					
P							
H	11.	.310	.330	.345	.310	.340	.327
A	12.	.340	.290	.325	.350	.315	.324
S	13.	.300	.290	.310	.300	.375	.315
E	14.	.350	.315	.355	.340	.350	.342
	15.	.370	.295	.375	.345	.370	.351
2							
		1 Block = 100 Trials in Phase 2					
P							
H	16.	.350	.350	.360	.330	.260	.330
A	17.	.340	.305	.360	.390	.335	.346
S	18.	.340	.365	.320	.275	.230	.306
E	19.	.310	.295	.300	.340	.300	.309
	20.	.230	.290	.320	.335	.320	.299
3							
		1 Block = 100 Trials in Phase 3					

Table 2: Prediction accuracy (by block and subject), Experiment 1.

	Phase 1	Phase 2	Phase 3
Subject:			
1	.400	.334	.314
2	.362	.304	.321
3	.365	.342	.332
4	.366	.329	.334
5	.370	.350	.289
Mean:	.373	.332	.318

Table 3: Average prediction accuracy (by phase and subject), Experiment 1.

Position in sequence:	Probability of correct response (N = 5):
1	.875
2	.707
3	.615
4	.574
5	.558
6	.517
7	.542
8	.639
9	.892

Table 4: Memory for serial position, Experiment 1.

Subject	Prediction accuracy in first 500 trials of Phase I (with memory question)	Prediction accuracy in first 100 trials after memory question discontinued
1	.385	.398
2	.355	.350
3	.370	.370
4	.330	.325
5	.335	.330
Mean:	.355	.355

Table 5: Comparison of prediction accuracy (by subject) with and without a memory question, Experiment 1.

If subjects are using one of the two crucial events exclusively to predict the 10th event's location, they should respond according to these rules:

When the 3rd event is in:	Then predict Box #:	I With this frequency:	II Actual prediction frequency (N = 5):
Box 1	1	0.50	0.374
	2	0.50	0.357
	3	0	0.268
Box 2	1	0.33	0.392
	2	0.33	0.340
	3	0.33	0.268
Box 3	1	0.50	0.378
	2	0	0.321
	3	0.50	0.302

When the 7th event is in:	Then predict Box #:	With this frequency:	Actual prediction frequency (N = 5):
Box 1	1	0.50	0.310
	2	0	0.340
	3	0.50	0.350
Box 2	1	0.33	0.389
	2	0.33	0.324
	3	0.33	0.287
Box 3	1	0.50	0.444
	2	0.50	0.368
	3	0	0.188

Table 6: Comparison of predicted and observed prediction probabilities, Experiment 1.

LOCATION OF SPOT ON SECOND APPEARANCE:	LOCATION OF SPOT ON FOURTH APPEARANCE:	LOCATION OF SPOT ON SIXTH APPEARANCE: I DURING TRAINING ([illegible],430 TRIALS):	II DURING SHIFTED FEEDBACK CONDITION ([illegible] TRIALS):	III DURING RANDOM FEEDBACK CONDITION ([illegible] TRIALS):
BOX 1	BOX 1			
BOX 2	BOX 2	BOX 1	BOX 2	BOX 1, 2, OR [illegible] p = 0.[illegible]
BOX 3	BOX 3			
BOX 1	BOX 2			
BOX 2	BOX 3	BOX 2	BOX 3	BOX 1, 2, OR [illegible] p = 0.[illegible]
BOX 3	BOX 1	(Clockwise movements)		
BOX 1	BOX 3			
BOX 3	BOX 2	BOX 3	BOX 1	BOX 1, 2, OR [illegible] p = 0.[illegible]
BOX 2	BOX 1	(Counter-clockwise movements)		

Table 7: Summary of rules governing feedback in Experiment 2.

	Block	Subject: 1	2	3	4	5	6	Mean
	1.	.342	.391	.333	.379	.374	.325	.357
P	2.	.383	.383	.379	.370	.292	.341	.358
H	3.	.374	.407	.354	.366	.379	.353	.372
A	4.	.399	.403	.366	.416	.366	.387	.390
S	5.	.374	.514	.379	.416	.399	.379	.410
E	6.	.395	.506	.391	.449	.395	.416	.425
	7.	.391	.457	.432	.436	.444	.416	.429
1	8.	.397	.436	.416	.403	.457	.412	.420
	9.	.407	.481	.420	.416	.453	.428	.434
	10.	.407	.461	.498	.461	.428	.444	.450
P								
H	11.	.337	.395	.313	.296	.321	.338	.333
A	12.	.391	.420	.366	.362	.399	.362	.383
S	13.	.465	.436	.481	.387	.412	.379	.427
E	14.	.379	.412	.420	.436	.465	.395	.421
2								
P								
H	15.	.280	.309	.313	.383	.296	.255	.306
A	16.	.337	.383	.325	.333	.272	.325	.329
S	17.	.428	.383	.383	.383	.313	.317	.366
E	18.	.366	.354	.370	.428	.354	.267	.357
3								

Note: All blocks = 243 trials.

Table 8: Prediction accuracy (by block and subject), Experiment 2.

Subject:	Phase 1	Phase 2	Phase 3
1	.387	.398	.353
2	.444	.416	.357
3	.397	.395	.348
4	.411	.370	.382
5	.399	.399	.309
6	.390	.369	.291
Mean:	.405	.391	.340

Table 9: Average prediction accuracy (by phase and subject), Experiment 2.

PHASE 1

	Pattern	Correct Response	Probability of responding correctly by subject: 1	2	3	4	5	6	Mean	[illegible]
Continuous patterns	11111	1	.3	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	
	22222	1	.5	[illegible]	.5	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
	33333	1	.6	[illegible]	.5	[illegible]	[illegible]	.7	[illegible]	
Single alternating patterns	12121	1	.4	[illegible]	0	[illegible]	.3	.3	[illegible]	
	21212	1	.4	.9	.1	.4	.3	.7	.467	
	23232	1	.5	.3	.7	[illegible]	[illegible]	.1	[illegible]	[illegible]
	32323	1	.3	.6	.5	.6	[illegible]	.3	[illegible]	
	13131	1	.4	.5	.6	.5	.3	.4	[illegible]	
	31313	1	.7	.6	.5	.3	.5	.6	[illegible]	
Clockwise patterns	12312	3	.6	.3	.1	.6	.4	.6	[illegible]	
	23123	3	.2	.2	.2	.1	.4	.2	[illegible]	[illegible]
	31231	3	.3	.1	0	[illegible]	.3	.2	[illegible]	
Counter-clockwise patterns	32132	2	0	.1	0	.1	.2	.4	[illegible]	
	21321	2	.3	.2	.2	.2	.3	.3	[illegible]	[illegible]
	13213	2	.4	.3	.7	.7	[illegible]	.7	[illegible]	
Double alternating patterns	11221	2	.3	.4	.5	.2	.4	.2	[illegible]	
	22112	3	.1	.1	.3	.3	.3	.2	[illegible]	
	22332	2	.8	.6	.1	.5	.2	.5	[illegible]	
	33223	3	.1	0	.3	.2	0	.1	[illegible]	[illegible]
	11331	3	.4	.1	.3	.2	.2	.2	.233	
	33113	2	.2	.5	.1	.2	.2	.3	.250	
Triple alternating patterns	11122	2	.7	.7	0	.3	.4	.6	.450	
	22211	3	.1	.5	.7	.6	.5	.3	.450	
	11133	3	.4	.4	.1	.3	.2	.4	.300	[illegible]
	33311	2	0	.1	.4	.3	.2	.3	.217	
	22233	2	.3	.4	.5	.3	.2	.3	.333	
	33322	3	.2	.1	.1	.3	0	.3	.167	

Table 10: Probability of responding correctly to salient patterns (by subject), Experiment 1, Phase I.

Subject #	Overall Phase I accuracy	Overall Phase I accuracy in the 2340 trials without the 9 highly salient patterns
1	.387	.382
2	.444	.440
3	.397	.394
4	.411	.410
5	.399	.395
6	.390	.385
Mean:	.405	.401

Table 11: Comparison of accuracy (by subject) with and without the 9 highly salient patterns, Experiment 2.

Pattern:	Block: 1	2	3	4	5	6	7	8	9	10
11111										
1	C	I	C	C	C	C	C	C	C	C
2	I	I	I	I	C	C	C	C	C	C
3	C	C	I	C	I	I	C	C	C	C
4	I	I	C	C	C	C	I	C	C	C
5	C	I	C	I	C	C	I	I	C	I
6	C	C	C	I	I	C	C	I	I	C
22222										
1	I	C	I	I	I	I	C	C	C	C
2	C	I	I	C	I	C	C	C	C	C
3	C	I	I	I	I	C	C	I	C	C
4	I	I	I	I	I	I	I	I	I	I
5	I	C	C	C	C	C	C	C	C	C
6	C	C	I	C	C	C	C	C	C	C
33333										
1	C	I	C	I	C	C	I	C	C	I
2	I	I	I	C	I	I	C	I	C	C
3	I	C	I	I	I	C	C	I	C	C
4	C	I	I	C	C	I	I	I	C	C
5	I	C	C	C	C	C	C	C	C	C
6	C	C	I	C	C	I	C	C	C	I
12121										
1	I	I	I	I	C	I	I	C	C	C
2	C	I	I	C	I	I	C	I	I	C
3	I	I	I	I	I	I	I	I	I	I
4	C	I	I	I	I	I	C	C	C	C
5	I	C	I	I	I	C	C	I	I	I
6	C	C	I	I	I	I	C	I	I	I
21212										
1	I	C	I	I	C	C	I	I	I	C
2	C	I	C	C	C	C	C	C	C	C
3	I	I	I	C	I	I	I	I	I	I
4	I	C	I	C	I	I	I	C	C	I
5	C	C	I	I	I	I	I	C	I	I
6	C	C	I	C	C	I	I	C	C	C

Table 12: Response accuracy (by block and subject) for each of the 9 highly salient patterns, Experiment 2, Phase I.

Pattern:	Block:									
	1	2	3	4	5	6	7	8	9	10
23232										
1	C	I	I	I	I	I	C	C	C	C
2	I	C	I	C	C	I	I	I	I	I
3	I	C	C	C	C	C	I	I	C	C
4	I	C	C	I	C	I	C	C	I	I
5	C	C	I	I	I	I	C	C	C	I
6	I	I	I	I	I	I	I	C	I	I
32323										
1	I	I	I	I	C	C	I	C	I	I
2	C	I	C	I	C	C	I	C	C	I
3	C	C	C	C	I	C	I	I	I	I
4	I	I	C	I	C	I	C	C	C	C
5	I	C	I	I	I	I	I	I	C	C
6	I	I	C	C	I	I	I	I	I	C
13131										
1	I	I	I	I	I	C	C	C	C	I
2	I	I	C	C	I	I	I	C	C	C
3	C	C	I	I	I	C	C	C	C	I
4	I	C	I	C	I	C	I	C	I	C
5	I	I	I	C	I	I	I	I	C	C
6	I	I	I	I	I	I	C	C	C	C
31313										
1	I	I	C	I	C	C	C	C	C	C
2	C	I	I	I	C	I	C	C	C	C
3	I	I	I	C	I	C	C	C	C	I
4	C	C	I	I	C	I	I	I	I	I
5	C	C	C	C	I	C	I	I	I	I
6	C	C	C	I	I	I	C	C	I	C

Note: 'C' and 'I' refer to correct and incorrect, respectively.

Table 12 (continued).

Continuous patterns

Block:	1	2	3	4	5	6	7	8	9	10
Number correct:	10	8	7	10	10	12	13	11	16	14
P(correct response):	.56	.44	.39	.56	.56	.67	.72	.61	.89	.78
Average of underlined pair of blocks:	.50		.48		.62		.66		.84	

Total correct responses to continuous patterns = 111
P(correct response to continuous patterns) = 0.617

Single-alternating patterns

Block:	1	2	3	4	5	6	7	8	9	10
Number correct:	15	17	11	14	13	13	16	22	19	18
P(correct response):	.42	.47	.30	.39	.36	.36	.44	.61	.53	.50
Average of underlined pair of blocks:	.44		.34		.36		.52		.52	

Total correct responses to continuous patterns = 158
P(correct response to continuous patterns) = 0.439

Table 13: Response accuracy (by block) for each of the 2 types of highly salient patterns, Experiment 2, Phase I.

	Subject:					
	1	2	3	4	5	6
Prediction accuracy in the last 81 trials of Phase I:	.375	.405	.410	.440	.480	.400
Rank:	6th	4th	3rd	2nd	1st	5th
Prediction accuracy in the first 81 trials of Phase II:	.250	.275	.300	.310	.300	.230
Rank:	4th	3rd	2nd	1st	2nd	5th
	(Note tie between subjects 3 and 5)					
Prediction accuracy in the first 81 trials of Phase II if Phase I feedback rules were still in effect:	.380	.400	.395	.450	.470	.400
Rank:	5th	3rd	4th	2nd	1st	3rd
	(Note tie between subjects 2 and 6)					

Table 14: Comparison of prediction accuracy (by subject) at the boundary of Phases I and II, Experiment 2.

	Subject: 1	2	3	4	5	6
Prediction accuracy in the last 81 trials of Phase II:	.481	.444	.469	.420	.395	.531
Rank:	2nd	4th	3rd	5th	6th	1st
Prediction accuracy in the first 81 trials of Phase III:	.325	.378	.310	.300	.338	.321
Rank:	3rd	1st	5th	6th	2nd	4th
Prediction accuracy in the first 81 trials of Phase III if Phase II feedback rules were still in effect:	.480	.444	.465	.422	.398	.539
Rank:	2nd	4th	3rd	5th	6th	1st

Table 15: Comparison of prediction accuracy (by subject) at the boundary of Phases II and III, Experiment 2.

	Pattern:	Correct Response:
1.	12231	2
2.	22323	1
3.	32322	1
4.	21132	3
5.	33313	2
6.	13311	2
7.	21111	1
8.	32321	1
9.	32132	2
10.	12211	3
11.	22333	2
12.	21133	3
13.	31331	3
14.	12323	1
15.	32112	3
16.	23122	3
17.	33333	1
18.	31212	1
19.	23132	1
20.	13112	2
21.	21223	2
22.	21332	3
23.	11132	3
24.	21323	2
25.	31121	2
26.	23321	3
27.	22132	2
28.	32121	1
29.	31231	3
30.	23232	1

Table 16: Patterns used in post-experimental "predict and justify" task, Experiment 3.

Block	1(E)	2(N)	3(E)	4(N)	5(E)	6(N)	7(E)	8(N)
1.	.317	.337	.329	.354	.342	.333	.329	.387
2.	.317	.362	.325	.333	.337	.370	.263	.317
3.	.329	.395	.358	.358	.284	.362	.337	.362
4.	.379	.350	.317	.342	.346	.317	.325	.358
5.	.321	.350	.342	.354	.309	.325	.350	.292
6.	.337	.383	.362	.354	.276	.280	.383	.346
7.	.333	.337	.362	.329	.292	.309	.329	.391
8.	.370	.313	.333	.358	.321	.366	.346	.362
9.	.383	.337	.337	.346	.329	.325	.305	.370
10.	.329	.374	.399	.374	.379	.342	.317	.432
11.	.309	.337	.366	.350	.346	.305	.337	.358
12.	.276	.337	.412	.416	.346	.321	.403	.374
13.	.387	.412	.354	.370	.942	.337	.333	.403
14.	.292	.395	.403	.387	.963	.337	.399	.374
15.	.309	.416	.346	.362	.955	.370	.325	.387
16.	.337	.379	.412	.387		.379	.395	.412
17.	.309	.374	.379	.305		.313	.407	.395
18.	.333	.486	.366	.370		.346	.337	.432
19.	.350	.424	.358	.329		.329	.366	.395
20.	.358	.436	.358	.416		.234	.387	.379
21.	.337	.453	.395	.251		.337	.359	.395
22.	.346	.432	.399	.317		.342	.362	.432
23.	.321	.440	.416	.325		.317	.395	.453
24.	.317	.366	.370	.370		.342	.346	.395
25.	.370	.444	.539	.370		.305	.403	.379
26.	.383	.436	.975	.333		.325	.407	.416
27.	.342	.453	.996	.383		.296	.403	.428
28.	.284	.469		.374		.350	.383	.424
29.	.342	.477		.391		.325	.403	.412
30.	.333	.453		.420		.333	.461	.407
31.	.366	.486		.387		.333	.420	.444
32.	.379	.465		.379		.300	.469	.416
33.	.342	.461		.403		.337	.432	.399
34.	.288	.481		.366		.346	.461	.403
35.	.321	.473		.309		.317	.469	.416
36.	.358	.436		.374		.300	.510	.407
37.	.354	.473		.444		.325	.461	.416
38.	.346	.543		.444		.329	.477	.391
39.	.313	.436		.449		.280	.502	.366
40.	.342	.469		.481		.317	.510	.403
41.	.346	.481		.477		.317	.502	.399
42.	.300	.453		.514		.370	.510	.440
43.	.300	.486		.449		.276	.535	.403
44.	.333	.568		.527		.321	.564	.374
45.	.288	.486		.527		.296	.551	.362
46.	.300	.556		.461		.342	.551	.453
47.	.329	.543		.486		.333	.556	.444
48.	.288	.539		.473		.362	.556	.374
49.	.292	.461		.523		.333	.560	.395
50.	.296	.449		.531		.321	.543	.424

Table 17: Prediction accuracy (by subject), Experiment 3.

Block:	Neutrally instructed subjects (N = 4):	Explicitly instructed subjects (N = 4):	All subjects (N = 8):
1.	.353	.329	.341
2.	.346	.311	.32[illegible]
3.	.369	.327	.348
4.	.342	.342	.342
5.	.330	.331	.331
6.	.341	.340	.341
7.	.342	.329	.336
8.	.350	.343	.348
9.	.345	.339	.342
10.	.381	.356	.369
11.	.338	.340	.339
12.	.362	.359	.361
13.	.381	.517*	.449*
14.	.373	.514*	.444*
15.	.384	.484*	.434*
16.	.389	.536	.463
17.	.347	.524	.436
18.	.409	.509	.459
19.	.369	.519	.444
20.	.379	.526	.453
21.	.359	.523	.441
22.	.381	.527	.454
23.	.384	.533	.459
24.	.368	.508	.438
25.	.375	.578*	.477*
26.	.378	.691*	.535*
27.	.390	.685*	.538*
28.	.404	.667	.535
29.	.401	.686	.544
30.	.403	.699	.551
31.	.413	.697	.555
32.	.390	.712	.551
33.	.400	.694	.547
34.	.399	.687	.543
35.	.379	.698	.539
36.	.379	.717	.548
37.	.415	.704	.560
38.	.427	.706	.567
39.	.383	.704	.544
40.	.418	.713	.566
41.	.419	.712	.566
42.	.444	.703	.574
43.	.404	.709	.557
44.	.448	.724	.586
45.	.418	.710	.564
46.	.453	.713	.583
47.	.452	.721	.587
48.	.437	.711	.574
49.	.428	.713	.571
50.	.431	.710	.571

Note: Data with an asterisk include accuracy information from the session in which a subject discovered the rule. Subsequent data reflect the assumption that starting with block 16 and 28, respectively, subjects 5(E) and 3(E) would have achieved 100% accuracy on each block.

Table 18: Average prediction accuracy (by block and instructional group), Experiment 3.

Neutrally-instructed subjects:	Explicitly-instructed subjects:	Explicitly-instructed subjects excluding perfect sessions:
2(N): .434	1(E): .331	1(E): .331
4(N): .395	3(E): .686	3(E): .367
6(N): .328	5(E): .835	5(E): .326
8(N): .396	7(E): .422	7(E): .422
Mean: .388	.568	.362

Note: Only subjects 3(E) and 5(E) had perfect sessions, that is, sessions after having discovered the rule in which their accuracy would have been 100% had they continued predicting.

Table 19: Average overall prediction accuracy (by subject), Experiment 3.

Subject:

Pattern	1(E)	2(N)	4(N)	6(N)	7(E)	8(N)	Mean
11111	.38	.94	.94	.58	.90	.92	.78
22222	.32	.90	.98	.24	.82	.94	.70
33333	.44	.92	.96	.28	.94	.92	.74
12121	.44	.80	.50	.40	.38	.60	.52
21212	.28	.76	.82	.40	.96	.48	.62
23232	.42	.86	.38	.32	.34	.50	.47
32323	.34	.70	.48	.30	.36	.42	.43
13131	.36	.70	.46	.32	.28	.50	.44
31313	.50	.74	.82	.26	.84	.34	.58

Note: Since this analysis is used to assess the contribution of the accuracy on these patterns to subjects' overall accuracy, data from the two subjects who discovered the rule are not included. Subjects 1(E) and 6(N) are the non-learners.

Table 20: Probability of responding correctly to highly salient patterns (by subject), Experiment 3.

Pattern:	Neutrally instructed subjects (N = 4)	Explicitly instructed subjects (N = 2)	All subjects (N = 6)
11111	.85	.64	.78
22222	.77	.57	.70
33333	.77	.69	.74
12121	.58	.41	.52
21212	.62	.62	.62
23232	.52	.38	.47
32323	.48	.35	.43
13131	.50	.32	.44
31313	.54	.67	.58

Note: Since this analysis is used to assess the contribution of the accuracy on these patterns to subjects' overall accuracy, data from the two subjects who discovered the rule are not included.

Table 21: Average probability of responding correctly to highly salient patterns (by non-solving subject), Experiment 3.

Subject:	1(E)			2(N)			4(N)		
Subject response:	1	2	3	1	2	3	1	2	3
Pattern:	Number of times (out of 50) each response given:								
11111	19	16	15	47	1	2	47	1	2
22222	16	19	15	45	3	2	49	1	0
33333	22	21	7	46	1	3	48	2	0
12121	22	20	8	40	6	4	25	21	4
21212	14	24	12	38	4	8	41	6	3
23232	21	13	16	43	3	4	19	7	24
32323	17	19	14	35	11	4	24	22	4
13131	18	14	18	35	7	8	23	5	22
31313	25	14	11	37	9	4	41	4	5

Subject:	6(N)			7(E)			8(N)		
Subject response:	1	2	3	1	2	3	1	2	3
Pattern:	Number of times (out of 50) each response given:								
11111	29	10	11	45	3	2	46	0	4
22222	12	22	16	41	6	3	47	1	2
33333	14	11	25	47	0	3	46	2	2
21212	20	15	15	19	29	2	30	9	11
21212	20	17	13	48	2	0	24	12	14
23232	16	20	14	17	0	33	25	9	16
32323	15	19	16	18	29	3	21	13	16
13131	16	16	18	14	2	34	25	19	6
31313	13	16	21	42	6	2	13	22	11

Note: Since this analysis is used to assess the contribution of the accuracy on these patterns to subjects' overall accuracy, data from the two subjects who discovered the rule are not included.

Table 22: Analysis (by subject) of sensitivity to salient patterns, Experiment 3.

Each Roman numeral represents the data from 5 consecutive blocks of 243 trials; each pattern appears once in each block of 243 trials and 5 times in each Roman numeral group.

Pattern:

11111	I	II	III	IV	V	VI	VII	VIII	IX	X
1(E)	.6	.6	.4	.4	.4	.2	.2	.2	.8	0
2(N)	.6	.8	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
4(N)	.6	1.0	1.0	.8	1.0	1.0	1.0	1.0	1.0	1.0
6(N)	1.0	.6	.2	.6	1.0	.6	.6	.6	.4	.2
7(E)	.2	.8	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
8(N)	.2	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0

22222	I	II	III	IV	V	VI	VII	VIII	IX	X
1(E)	.4	.6	.2	.6	.4	.2	.2	.4	.2	0
2(N)	.2	.8	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
4(N)	.8	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
6(N)	0	.4	0	0	.4	.2	.4	.8	.2	0
7(E)	.2	.4	.8	1.0	1.0	.8	1.0	1.0	1.0	1.0
8(N)	.4	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0

33333	I	II	III	IV	V	VI	VII	VIII	IX	X
1(E)	.4	.4	.4	.8	.4	.6	.4	.4	.2	.4
2(N)	.2	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
4(N)	.8	1.0	1.0	.8	1.0	1.0	1.0	1.0	1.0	1.0
6(N)	0	.6	0	.4	0	.2	.4	.2	.6	.4
7(E)	.4	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
8(N)	.4	.8	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0

12121	I	II	III	IV	V	VI	VII	VIII	IX	X
1(E)	.8	.6	.4	.4	.6	.2	.2	.4	.8	0
2(N)	.4	.4	.6	1.0	1.0	1.0	.8	1.0	.8	1.0
4(N)	0	.2	.2	.2	.2	.6	1.0	1.0	.8	.8
6(N)	.4	.6	.4	.6	.6	.4	.2	.2	.2	.4
7(E)	.4	0	0	0	0	0	.4	1.0	1.0	1.0
8(N)	.6	.8	.8	.6	.6	.2	.8	.4	.8	.6

Table 23: Probability of a correct response to each of the 9 highly salient patterns (by subject), Experiment 3.

21212	I	II	III	IV	V	VI	VII	VIII	IX	X
1(E)	.2	.2	.2	.4	.4	.6	.2	.4	.2	0
2(N)	.4	.2	.4	.8	1.0	1.0	1.0	.8	1.0	1.0
4(N)	.6	1.0	.6	.8	.8	.8	.8	.8	1.0	1.0
6(N)	0	.8	.2	0	.8	.4	.6	.8	.2	.2
7(E)	.8	.8	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
8(N)	.6	.4	.2	.4	.2	.8	.6	.8	.6	.2

23232	I	II	III	IV	V	VI	VII	VIII	IX	X
1(E)	.6	0	.2	0	.8	.6	.8	0	.8	.4
2(N)	.8	.6	.4	.8	1.0	1.0	1.0	1.0	1.0	1.0
4(N)	0	.2	.2	.2	0	0	.6	.8	.8	1.0
6(N)	.2	.2	.4	0	.6	0	.2	.6	.4	.6
7(E)	.2	0	.2	0	0	0	0	1.0	1.0	1.0
8(N)	0	.2	.4	.4	.4	.8	.8	.8	.6	.6

32323	I	II	III	IV	V	VI	VII	VIII	IX	X
1(E)	.2	.4	0	.2	.4	.4	.6	.2	.2	.8
2(N)	.4	.6	.6	.8	.4	.8	1.0	1.0	.8	.6
4(N)	0	.2	.2	.2	0	.6	.8	1.0	1.0	.8
6(N)	.2	.2	.4	.2	.4	.2	.2	.4	.4	.4
7(E)	0	.2	0	0	.2	0	.2	1.0	1.0	1.0
8(N)	.4	0	.6	0	.6	.8	.8	.4	.4	.2

13131	I	II	III	IV	V	VI	VII	VIII	IX	X
1(E)	.4	.2	.6	.6	.6	.4	0	.2	.2	.2
2(N)	.4	.2	0	.8	1.0	1.0	1.0	1.0	1.0	.6
4(N)	.2	0	0	0	.4	.2	1.0	1.0	1.0	.8
6(N)	.8	.4	.4	.2	.6	0	0	.2	.4	.2
7(E)	0	0	0	0	0	0	0	.8	1.0	1.0
8(N)	.8	.6	.4	.2	.4	.6	0	.6	.6	.8

31313	I	II	III	IV	V	VI	VII	VIII	IX	X
1(E)	.2	.8	.6	.6	.2	.4	.6	.6	.8	.2
2(N)	.6	0	.2	.6	1.0	1.0	1.0	1.0	1.0	1.0
4(N)	.8	.8	.6	.6	.8	.8	.8	1.0	1.0	1.0
6(N)	.2	.2	.4	.2	.2	.4	.4	.2	0	.4
7(E)	.4	.6	.6	1.0	1.0	1.0	.8	1.0	1.0	1.0
8(N)	.4	.2	.4	.2	.2	.4	.4	.2	.6	.4

Table 23 (continued).

Neutrally instructed subjects:

Subject:	Overall prediction accuracy:	Overall prediction accuracy in the 11,700 trials without the 9 highly salient patterns:
2(N)	.434	.419
4(N)	.395	.385
6(N)	.328	.326
8(N)	.396	.387
Mean:	.388	.379
	Explicitly-instructed subjects:	
1(E)	.331	.328
7(E)	.422	.413
Mean:	.377	.371

Note: Since this analysis is used to assess the contribution of the accuracy on these patterns to subjects' overall accuracy, data from the two subjects who discovered the rule are not included.

Table 24: Comparison of prediction accuracy (by subject) with and without data from highly salient patterns, Experiment 3.

Subject			Correct response (feedback) 1	2	3	Mean
1(E)	Correct responses	X:	873.33	798.00	622.67	777.00
		s.d.:	1521.17	1412.50	990.50	
	Incorrect responses	X:	771.83	781.17	799.67	784.67
		s.d.:	1365.83	1408.83	1378.17	
2(N)	Correct responses	X:	961.50	1221.33	1215.33	1106.33
		s.d.:	722.67	785.33	853.17	
	Incorrect responses	X:	1424.83	1235.50	1239.50	1285.33
		s.d.:	878.50	816.67	786.17	
4(N)	Correct responses	X:	1467.50	2457.33	2617.50	2042.33
		s.d.:	1958.67	2424.17	2268.17	
	Incorrect responses	X:	2316.00	2414.50	2263.50	2329.33
		s.d.:	2295.50	2445.50	2346.67	
6(N)	Correct responses	X:	648.50	677.67	729.50	684.00
		s.d.:	841.17	891.33	1253.50	
	Incorrect responses	X:	746.33	696.33	671.50	703.67
		s.d.:	1281.00	1013.67	899.33	
7(E)	Correct responses	X:	1021.33	1402.33	1390.33	1237.67
		s.d.:	913.17	1254.33	1324.83	
	Incorrect responses	X:	1430.33	1348.83	1336.83	1365.83
		s.d.:	1237.17	1203.00	1152.83	
8(N)	Correct responses	X:	1011.33	1177.33	1172.33	1103.17
		s.d.:	779.67	861.83	899.17	
	Incorrect responses	X:	1278.50	1149.00	1146.50	1182.50
		s.d.:	1029.00	802.00	812.67	

Test for significant difference between means to respond correctly and incorrectly (*significant at alpha = 0.5)

Subject	t(12,148)	Subject	t(12,148)
1(E)	0.289	6(N)	0.966
2(N)	12.05*	7(E)	5.91*
4(N)	6.67*	8(N)	4.96*

Table 25: Statistics for response latency (in milliseconds) (by non-solving subject), Experiment 3.

N	Daily session:	Event number: 1	2	3	4	5
		Average rank assigned:				
4	1	2.75	4.25	3.25	3.25	1.5
4	2	2.25	4.25	3.75	3.5	1.25
4	3	2.75	4.5	3.0	3.25	1.5
4	4	2.5	4.5	3.0	2.5	1.75
3	5	4.0	2.33	4.0	3.0	1.67
3	6	5.0	3.33	2.67	1.67	2.33
3	7	3.33	2.33	4.0	2.33	3.0
3	8	2.33	3.67	2.33	3.67	3.0
2	9	3.5	3.0	4.0	3.5	1.0
2	10	2.0	3.5	3.0	5.0	1.5
2	11	2.0	2.0	4.5	4.5	2.0
2	12	2.0	3.5	2.5	4.0	3.0
2	13	3.0	3.0	2.0	3.5	3.5
2	14	2.0	3.0	2.5	4.5	3.0
2	15	2.0	3.0	2.5	4.5	3.0
2	16	2.0	3.5	2.0	4.0	3.5
2	17	3.0	2.5	2.5	4.0	3.0

Mean rankings by subject:

Event #	1(E)	3(E)	5(E)	7(E)	Mean
1	3.53	3.50	2.25	1.82	2.78
2	3.65	3.38	3.75	3.24	3.51
3	2.12	3.38	4.25	3.82	3.39
4	3.94	2.00	3.50	3.53	3.24
5	1.76	2.75	1.25	2.65	2.10

Note: Subjects 3(E) and 5(E) did not perform the ranking task following the session in which they discovered the rule.

Table 26: Ranking task results, Experiment 3.

	Pattern:	Correct Response:	1(E)	2(N)	4(N)	6(N)	7(E)	8(N)	P(C)
1.	12231	2	2	2	3	2	2	2	.33
2.	22323	1	1	1	1	2	2	2	.50
3.	32322	1	3	1	1	2	1	3	.50
4.	21132	3	1	3	2	1	2	3	.33
5.	33313	2	2	3	1	3	1	1	.17
6.	13311	2	3	1	2	2	1	2	.50
7.	21111	1	1	3	1	2	1	1	.67
8.	32321	1	2	1	1	2	1	2	.50
9.	32132	2	1	3	2	3	3	3	.17
10.	12211	3	1	2	2	1	1	2	0
11.	22333	2	2	2	2	3	2	1	.67
12.	21133	3	2	2	3	3	2	1	.33
13.	31331	3	2	1	3	3	1	2	.33
14.	12323	1	3	1	1	2	1	1	.67
15.	32112	3	1	2	1	1	2	2	0
16.	23122	3	3	3	3	3	3	3	1.0
17.	33333	1	1	1	1	3	1	1	.33
18.	31212	1	2	1	1	2	1	1	.67
19.	23132	1	3	3	1	2	3	2	.17
20.	13112	2	1	2	1	1	1	2	.33
21.	21223	2	2	1	3	3	2	1	.33
22.	21332	3	3	2	2	2	3	2	.33
23.	11132	3	2	3	3	2	3	2	.50
24.	21323	2	1	1	2	1	3	2	.33
25.	31121	2	3	1	1	2	1	3	.17
26.	23321	3	2	3	3	2	2	1	.33
27.	22132	2	1	3	2	2	2	3	.50
28.	32121	1	3	1	1	1	1	3	.67
29.	31231	3	2	2	3	3	1	2	.33
30.	23232	1	2	3	1	2	1	3	.33
		Number correct:	9	14	21	9	15	10	.43

Note: Subjects 1(E) and 6(N) were non-learners.

Table 27: Prediction accuracy (by subject) in post-experimental "predict and justify" task, Experiment 3.

4 Spots in Box 1	Correct Response	4 Spots in Box 2	Correct Response	4 Spots in Box 3	Correct Response
11112	1	22221	1	33331	1
11113	1	22223	1	33332	1
11121	2	22212	3	33313	2
11131	3	22232	2	33323	3
11211	1	22122	1	33133	1
11311	1	22322	1	33233	1
12111	3	21222	2	31333	3
13111	2	23222	3	32333	2
21111	1	12222	1	13333	1
31111	1	32222	1	23333	1

Table A-1: Additional potentially salient patterns ("quads").

In Experiment 2, Phase I

Subject #	Overall Phase I accuracy	Overall Phase I accuracy in the 2040 trials without the 39 patterns ofinterest
1	.387	.369
2	.444	.433
3	.397	.386
4	.411	.407
5	.399	.384
6	.390	.378
Mean:	.405	.393

In Experiment 3 (among 4 learning subjects)

Subject #	Overall accuracy	Overall accuracy in the 10,200 trials without the 39 patterns of interest
2(N)	.434	.406
4(N)	.396	.365
7(E)	.422	.391
8(E)	.396	.368
Mean:	.412	.383

Table A-2: Comparison of accuracy (by subject) with and without 39 salient and potentially salient patterns, Experiment 2, Phase I, and Experiment 3.

For Experiment 2, Phase I (N = 6):

RESPONSE	FEEDBACK 1	2	3	P(Resp)
1	1555	1363	1381	.351
2	863	1731	1602	.343
3	822	1406	1517	.305

Note: In this grid, Experiment 2, Phase I consists of 2430 - (39 trials x 10 blocks) = 2040 trials per subject (or 12,240 trials for all six subjects combined); there are 3240 trials that predict Box 1 and 4500 trials that predict Boxes 2 and 3 each.

For Experiment 3 (Among the 4 learners):

RESPONSE	FEEDBACK 1	2	3	P(Resp)
1	4977	4408	4439	.339
2	2971	5601	5531	.346
3	2852	4991	5030	.316

Note: In this grid, Experiment 3 consists of 12,150 - (39 trials x 50 blocks) = 10,200 trials per subject (or 40,800 trials for all four learning subjects combined); there are 10,800 trials that predict Box 1 and 15,000 trials that predict Boxes 2 and 3 each.

Table A-3: Response pattern analysis without 39 patterns, Experiment 2, Phase I, and Experiment 3.

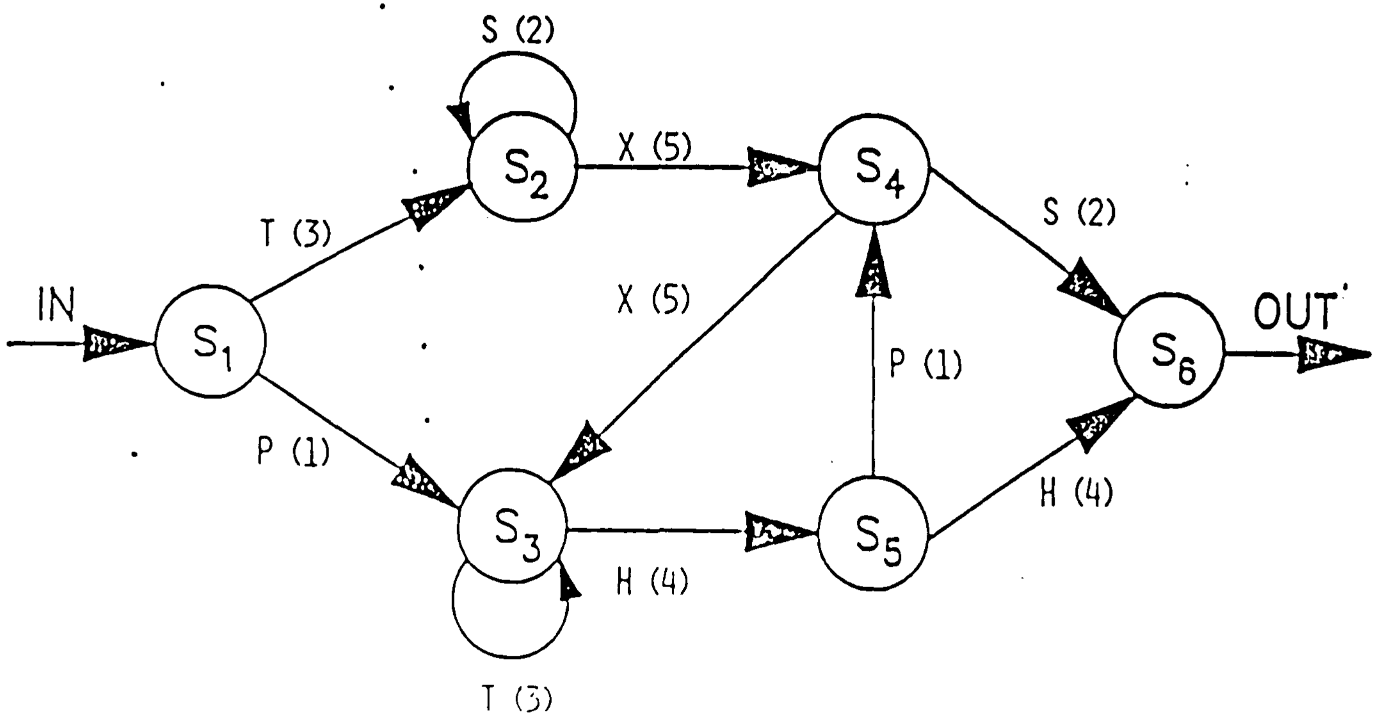

Figure 1: Example of a Markov system designed to generate strings of letters.

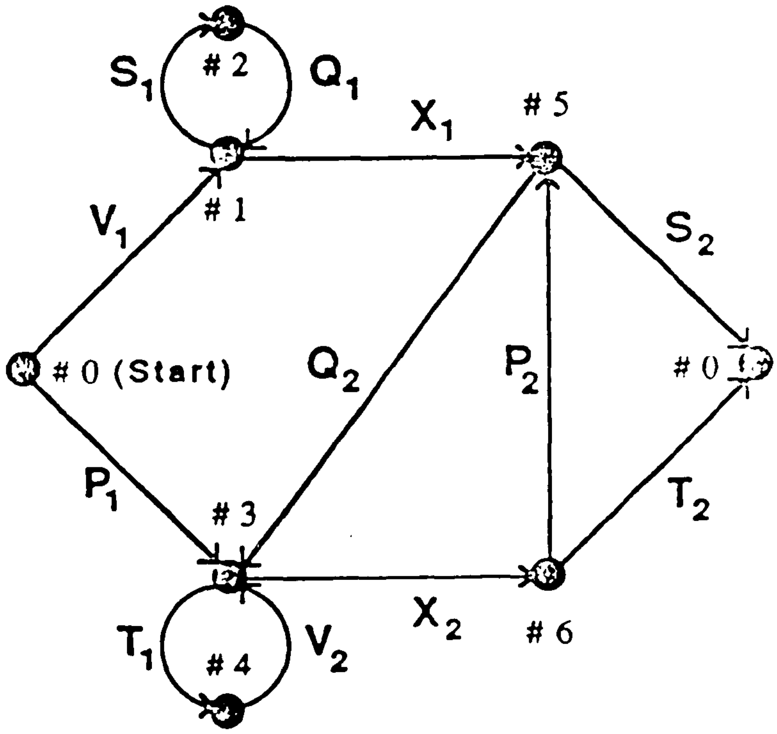

Figure 2: Markov system used by Cleeremans and McClelland (1991).

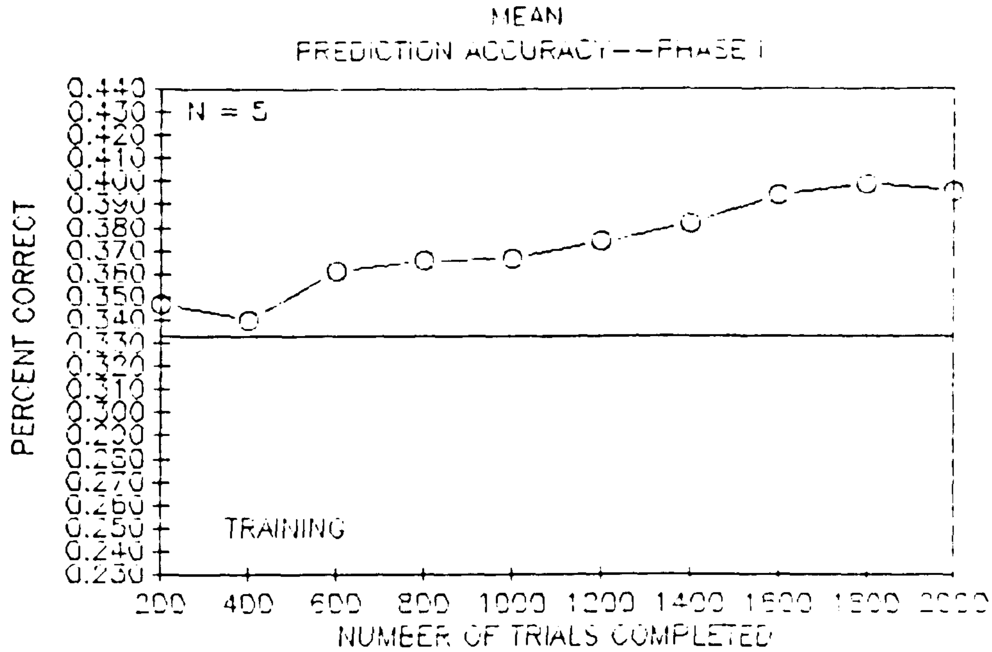

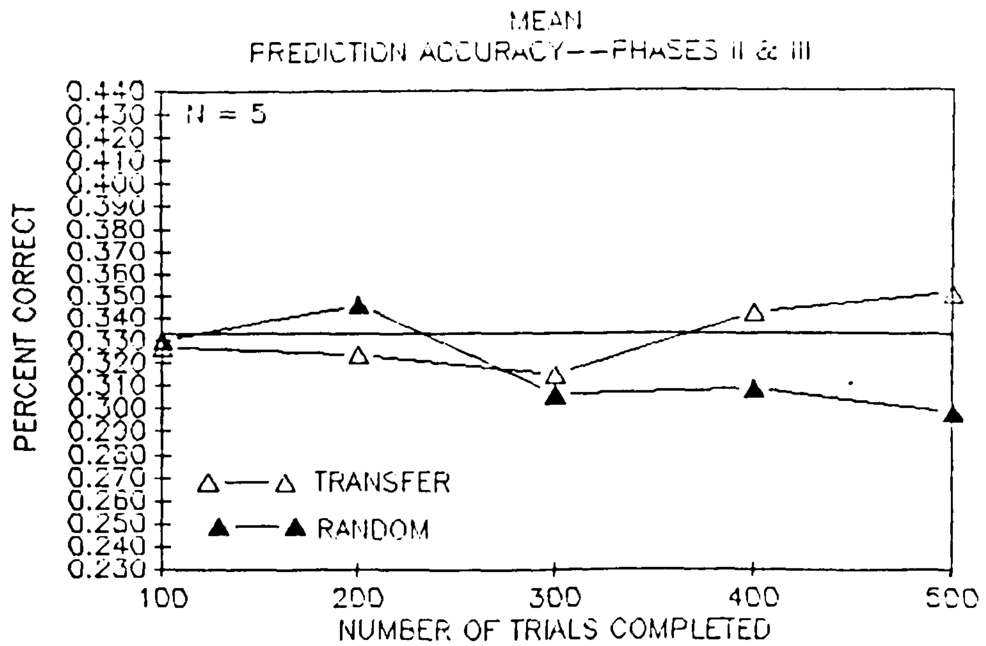

Figure 3: Average prediction accuracy, Experiment 1.

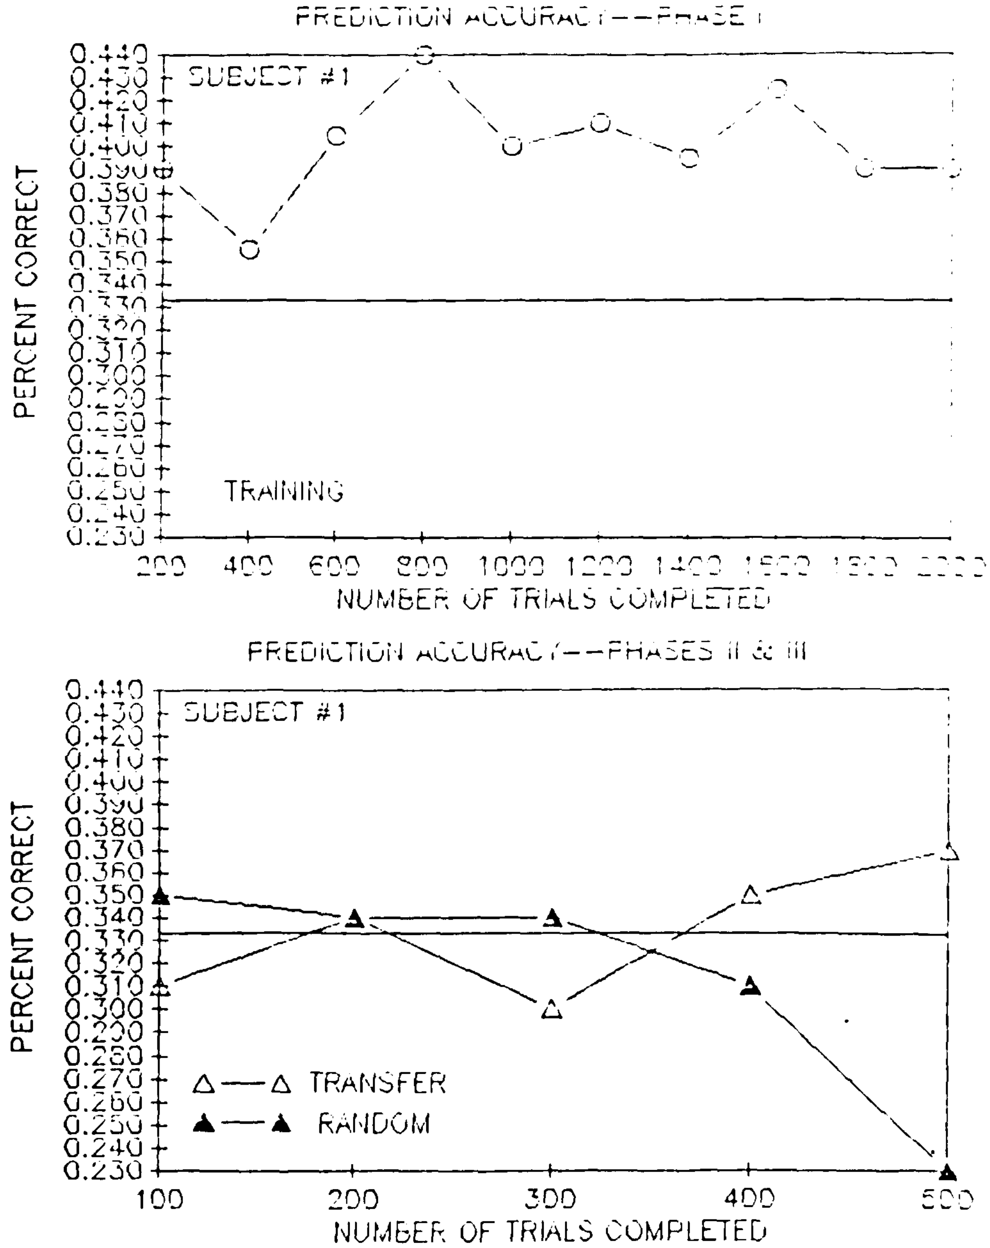

Figure 4: Prediction accuracy, Subject 1, Experiment 1.

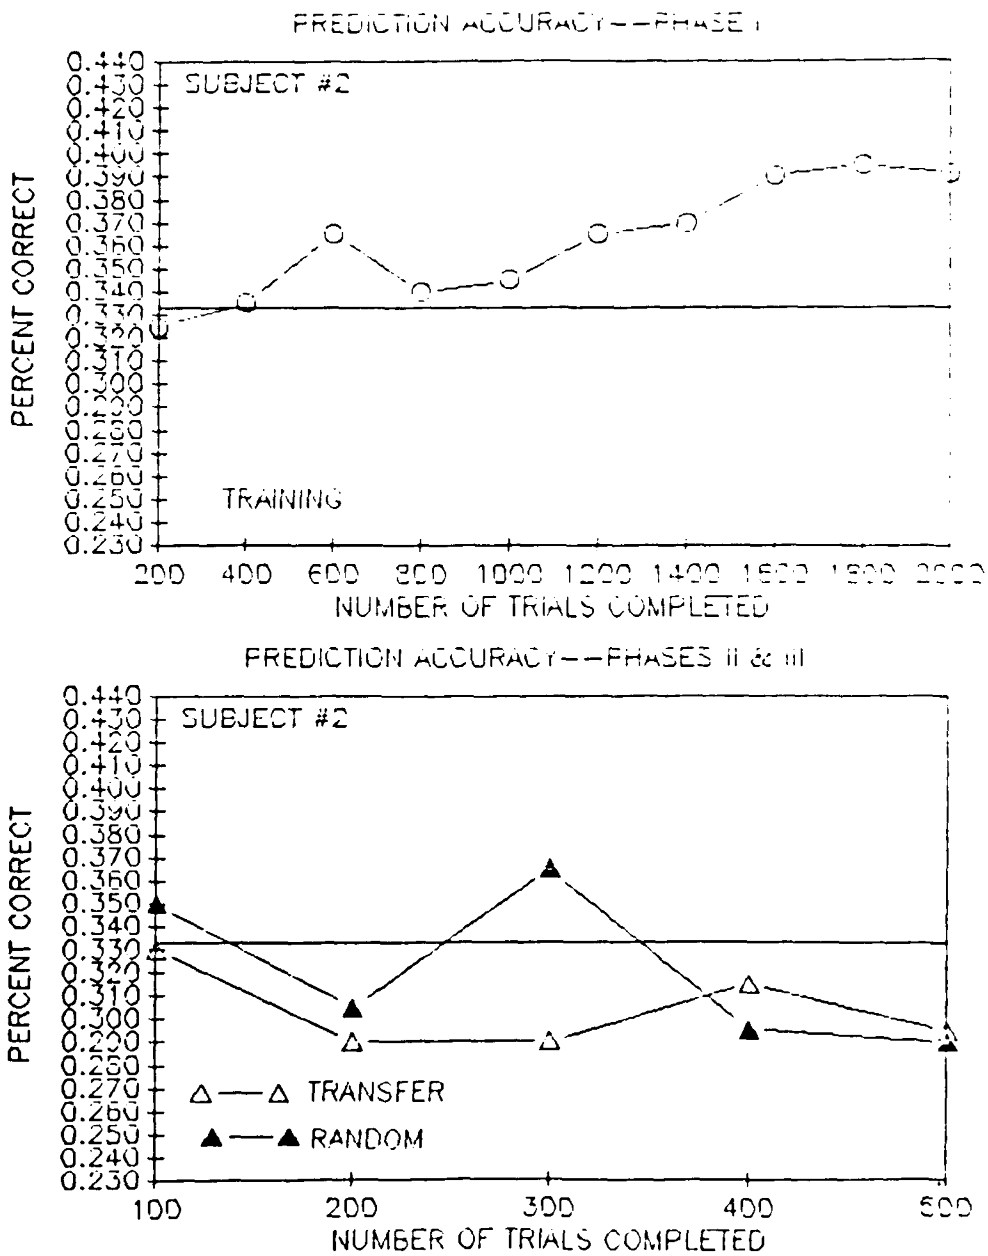

Figure 5: Prediction accuracy, Subject 2, Experiment 1.

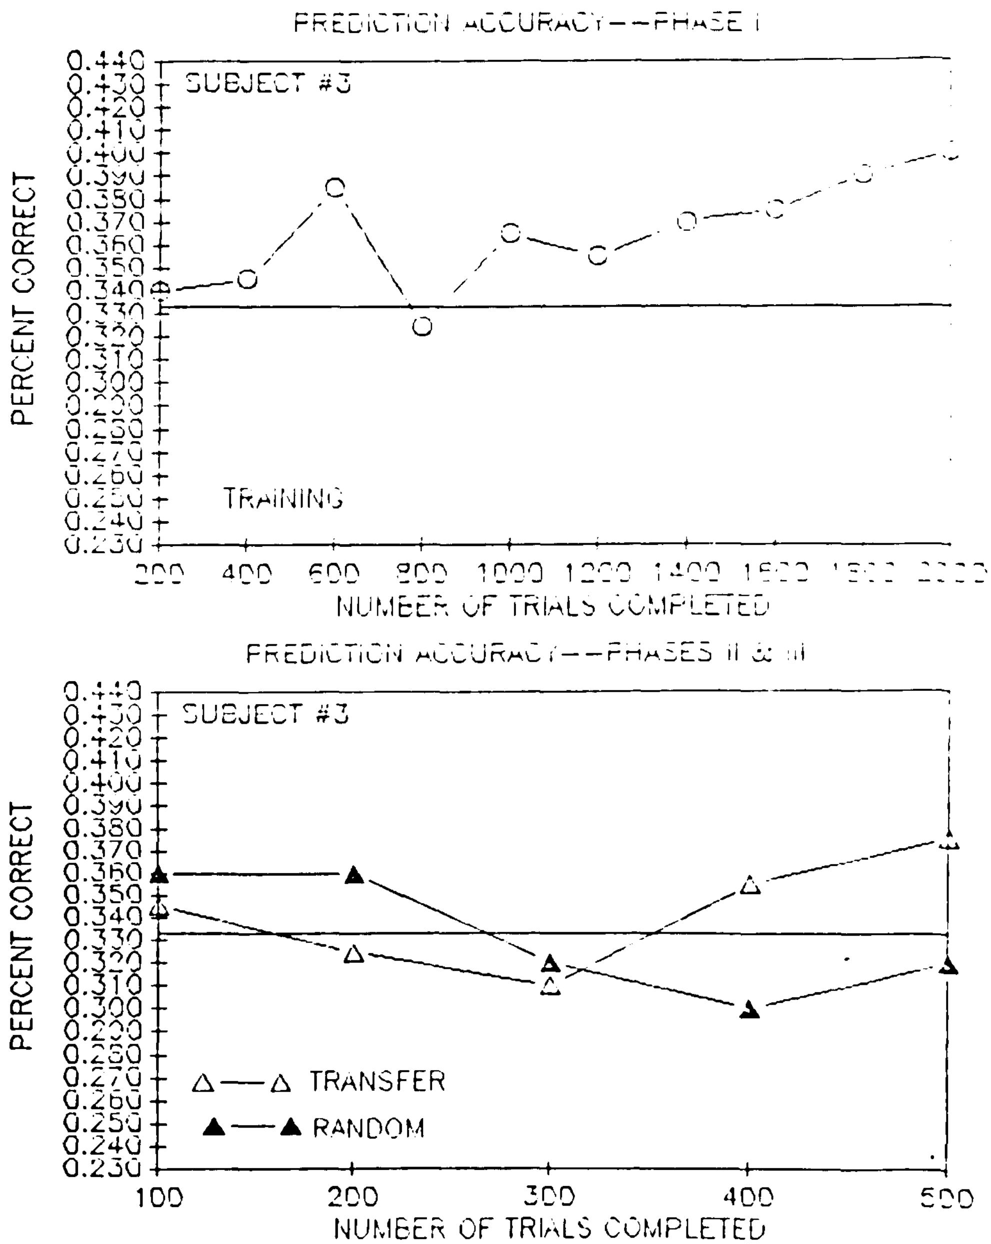

Figure 6: Prediction accuracy, Subject 3, Experiment 1.

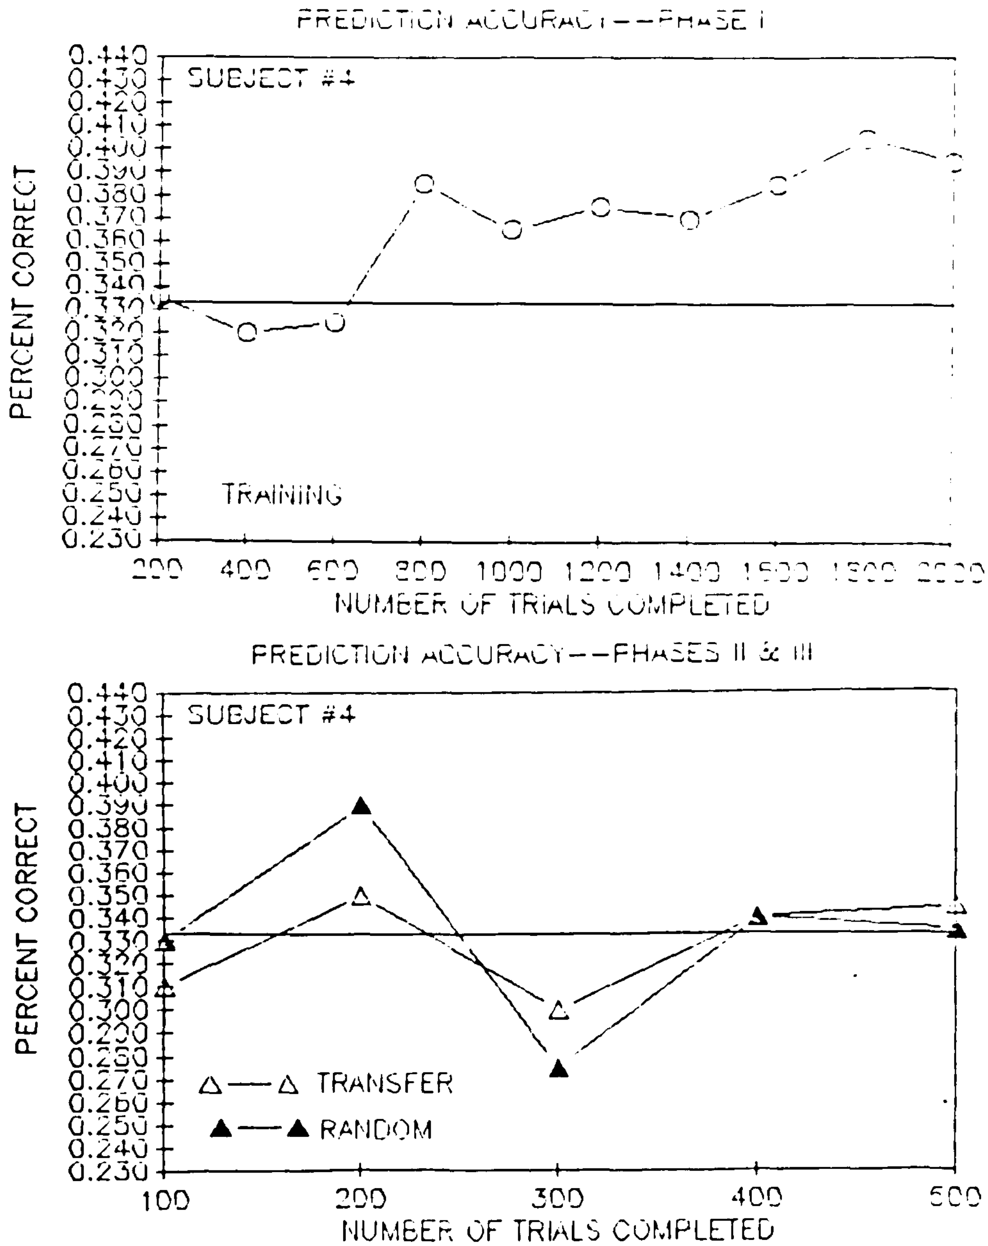

Figure 7: Prediction accuracy, Subject 4, Experiment 1.

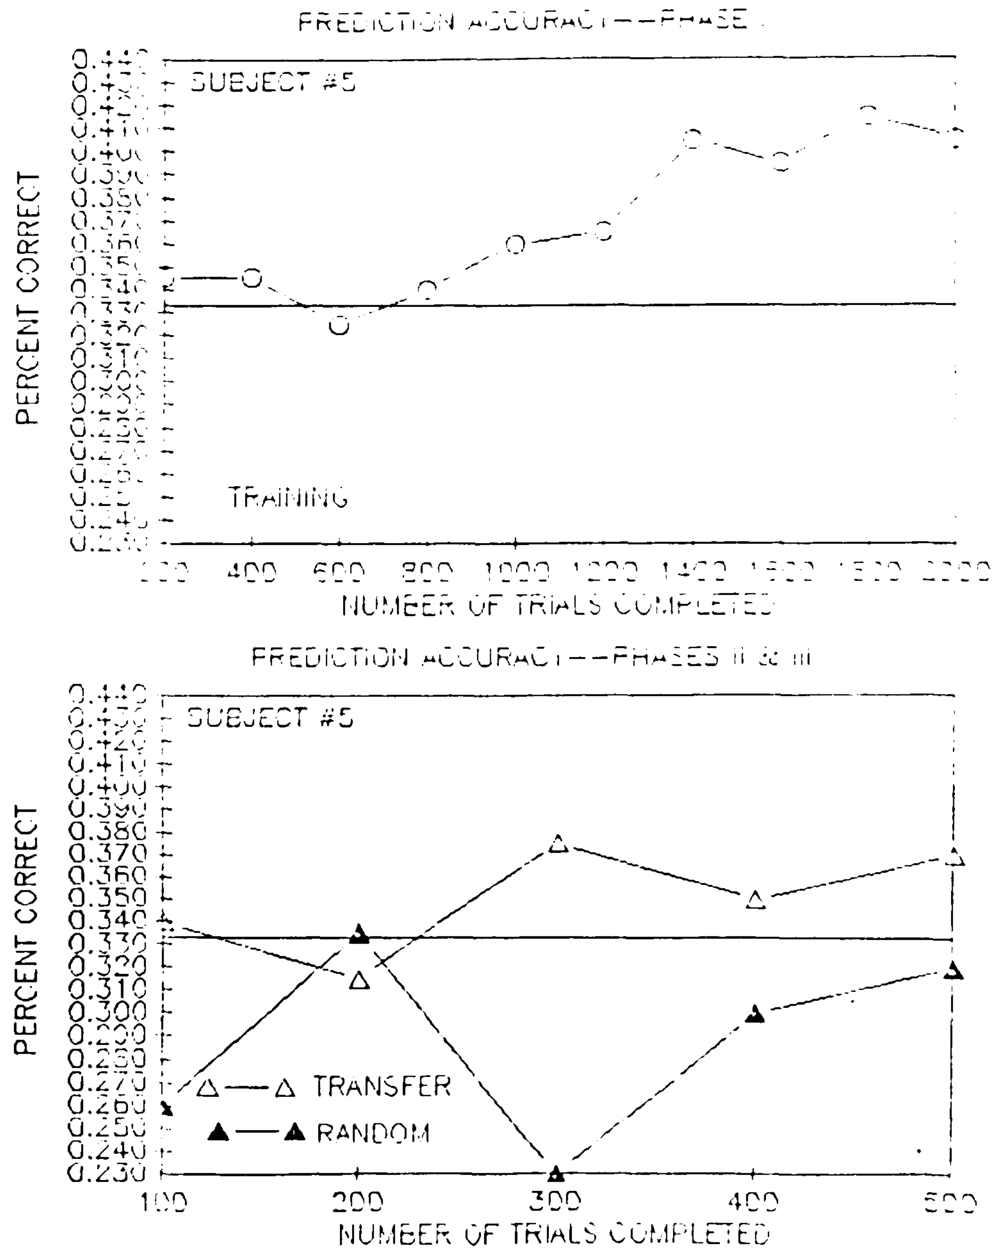

Figure 8: Prediction accuracy, Subject 5, Experiment 1.

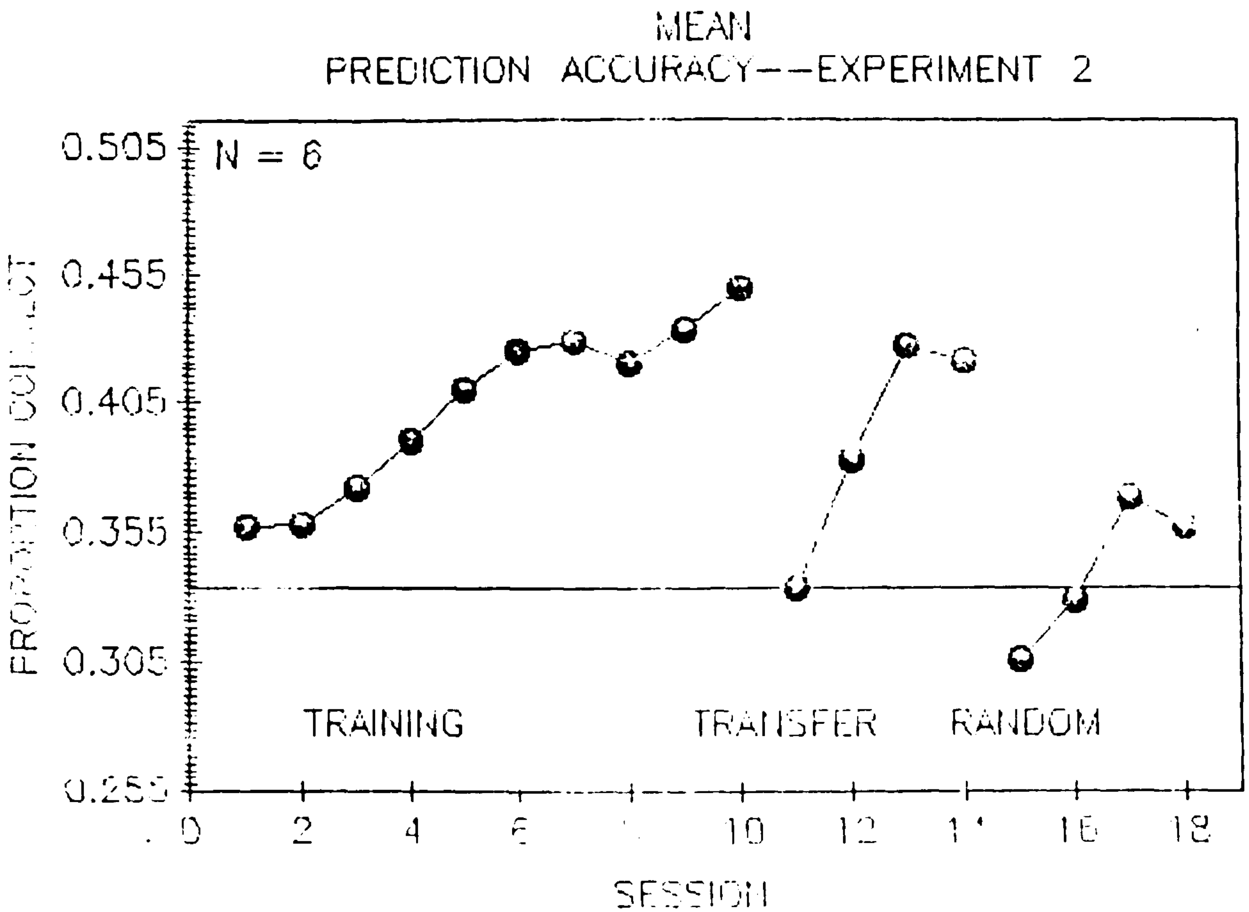

Figure [illegible]. Average prediction accuracy, Experiment 2.

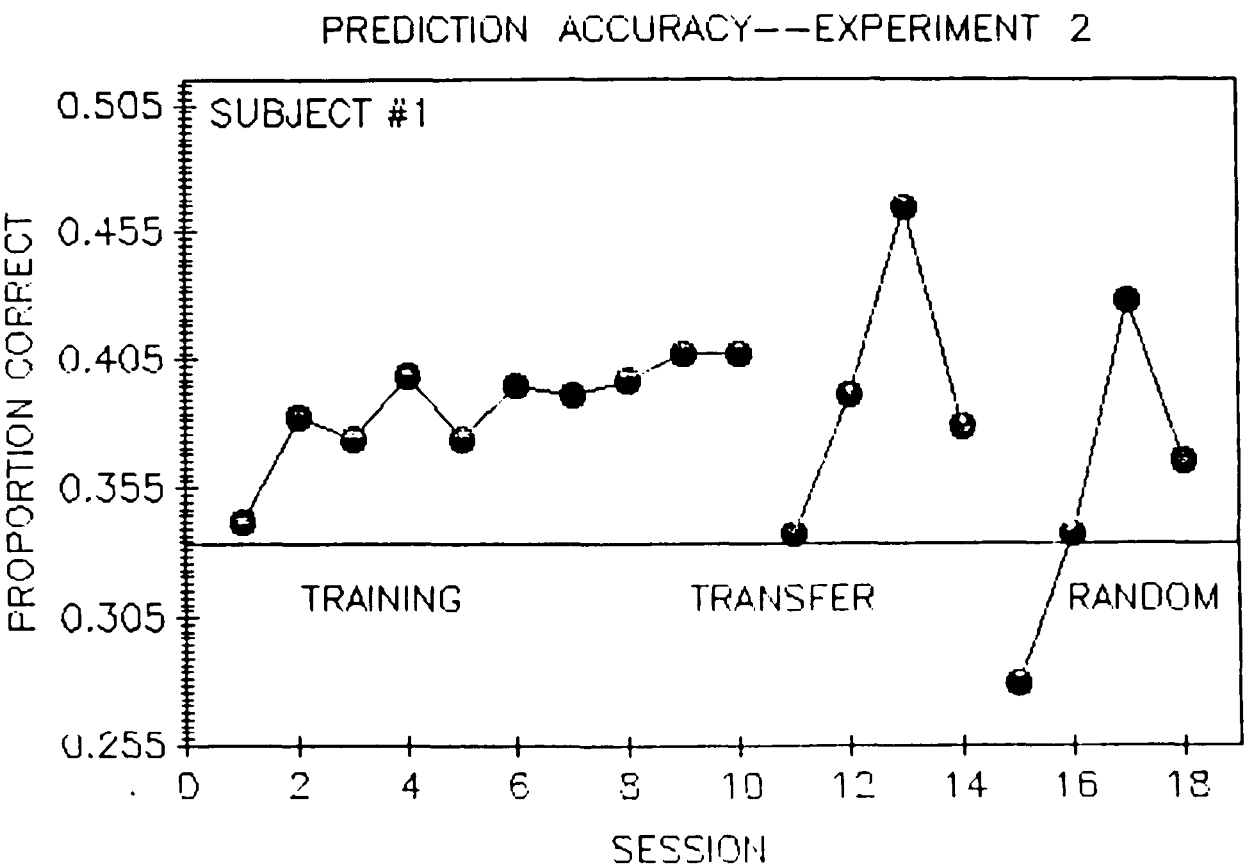

Figure 10: Prediction accuracy, subject 1, Experiment 2.

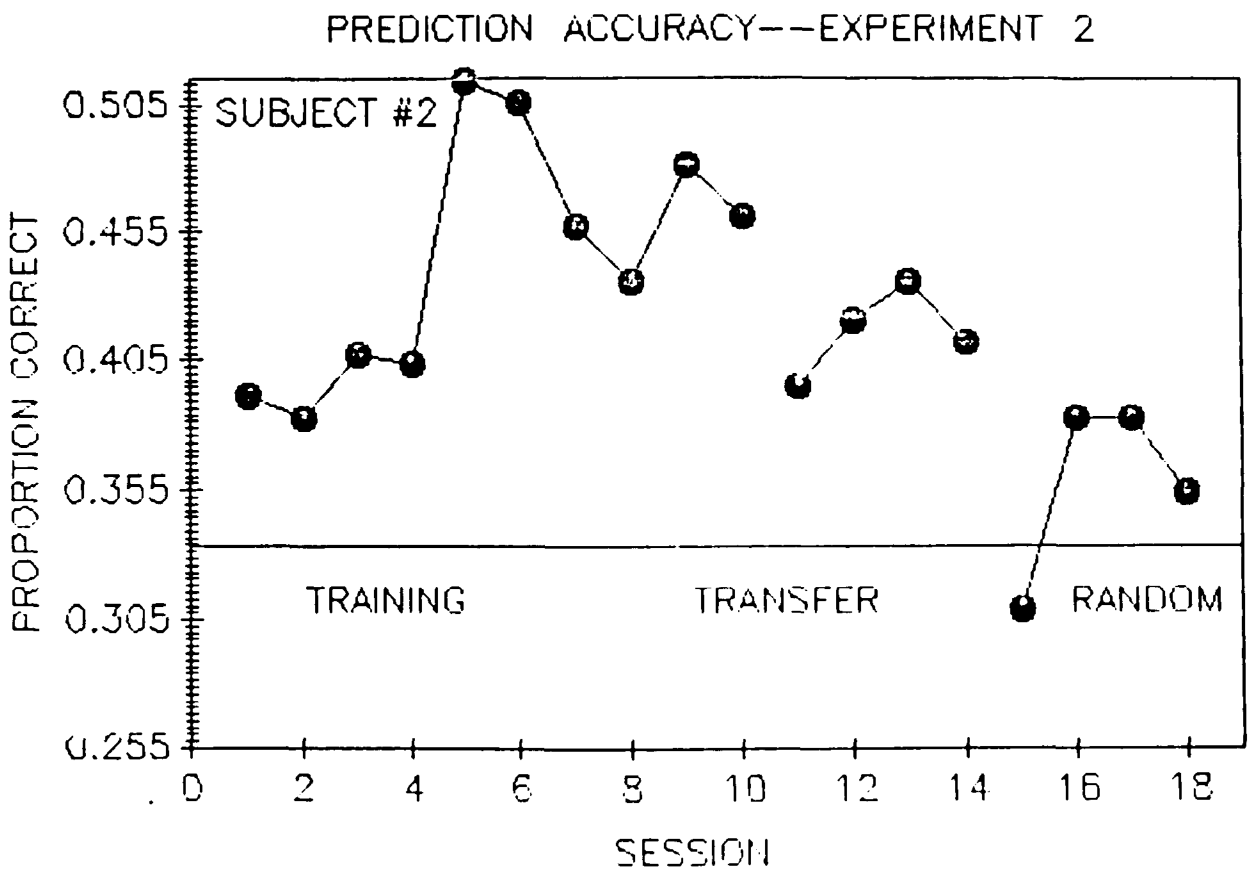

Figure 11: Prediction accuracy, Subject 2, Experiment 2.

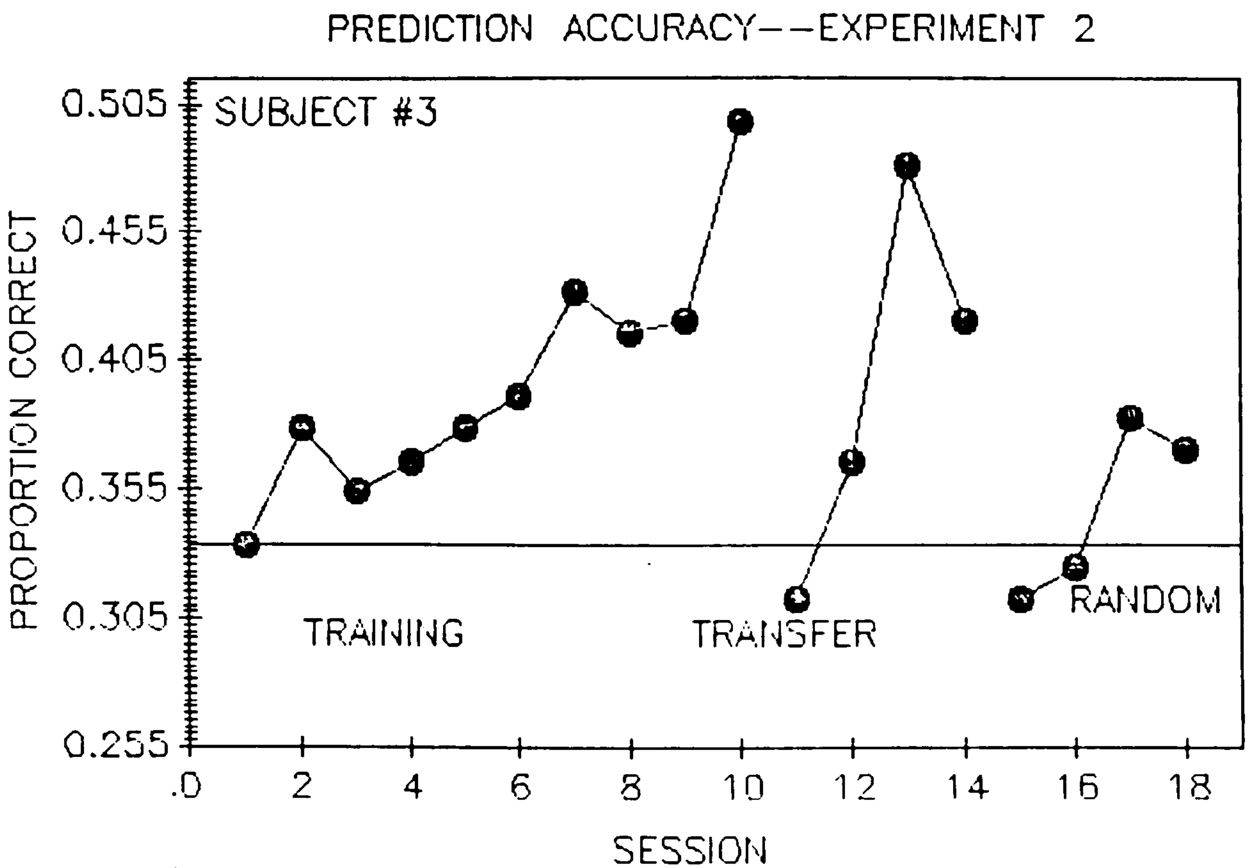

Figure 12: Prediction accuracy, subject 3, Experiment 2.

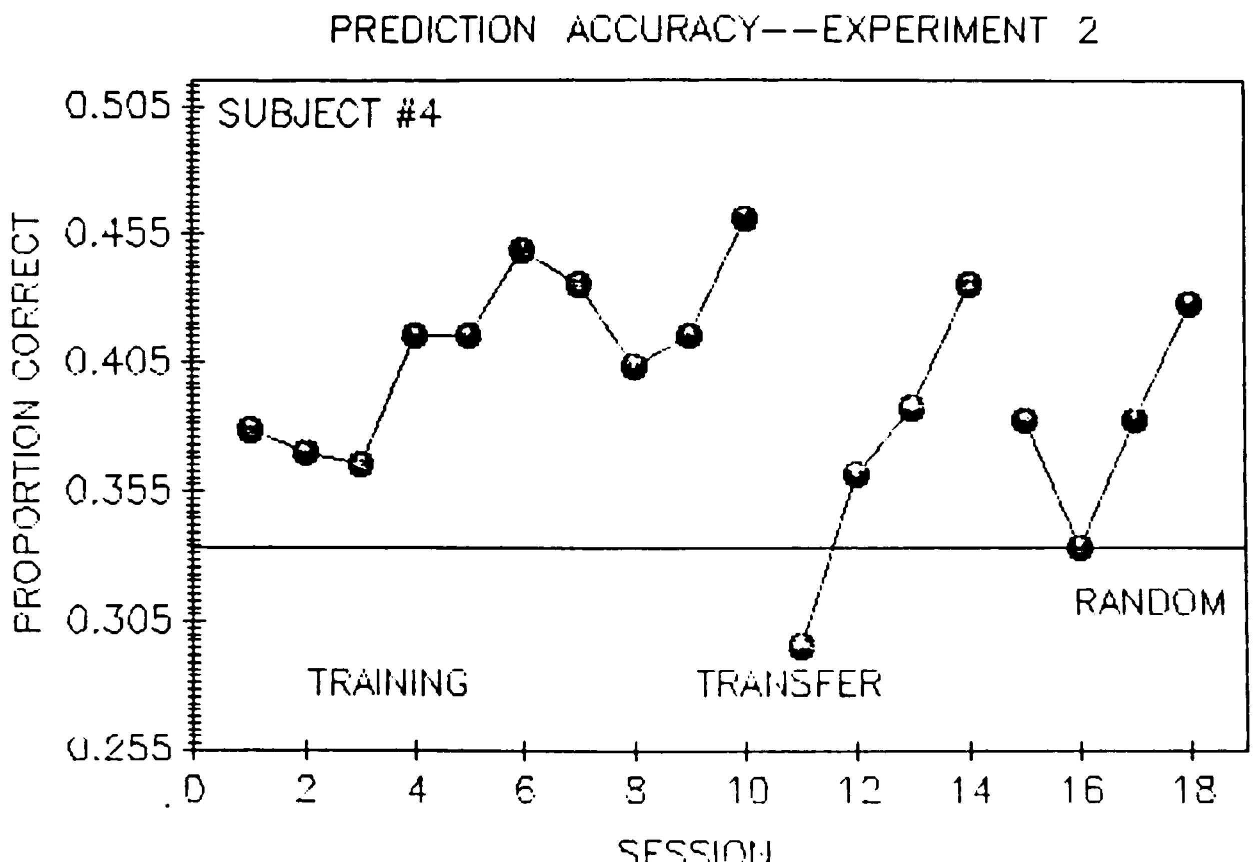

Figure 13: Prediction accuracy, subject 4, Experiment 2.

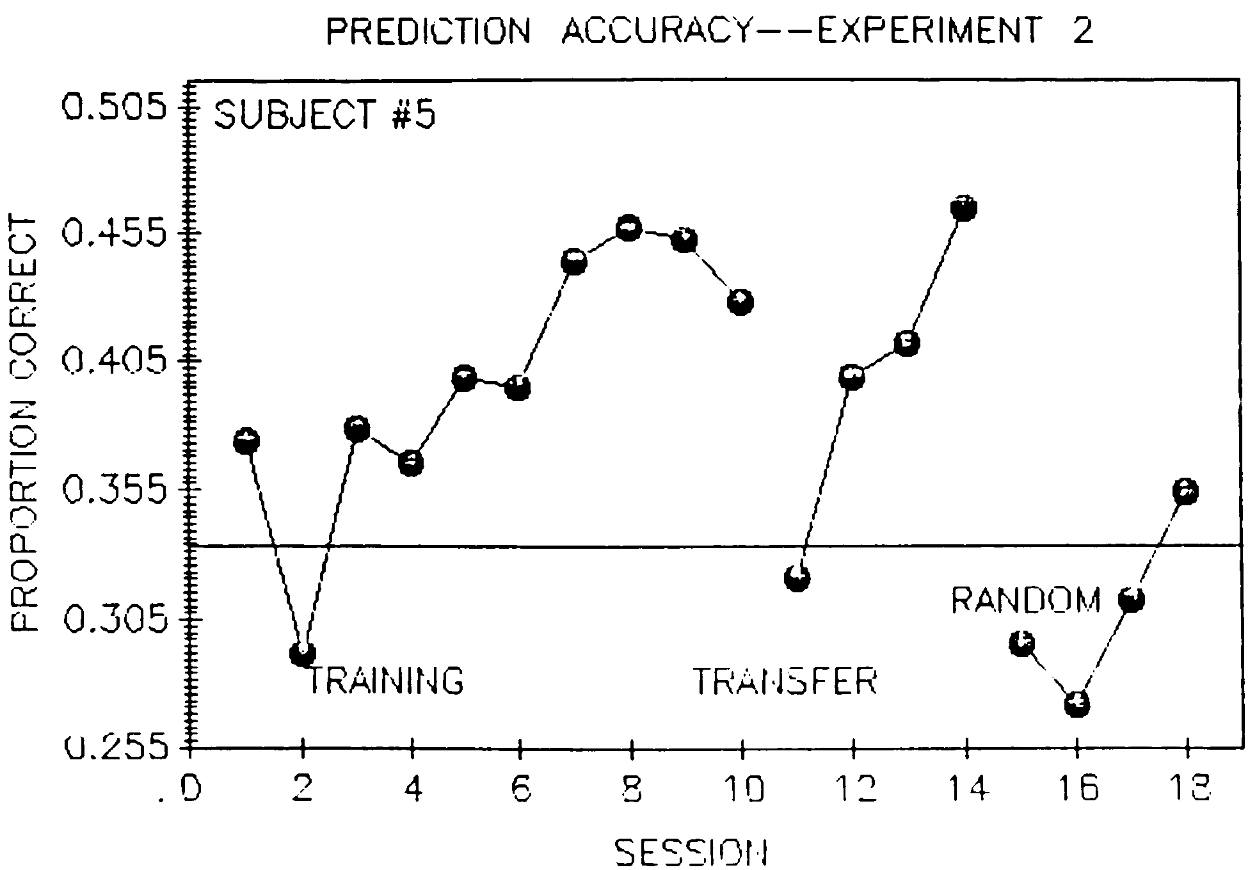

Figure 14: Prediction accuracy, subject 5, Experiment 2.

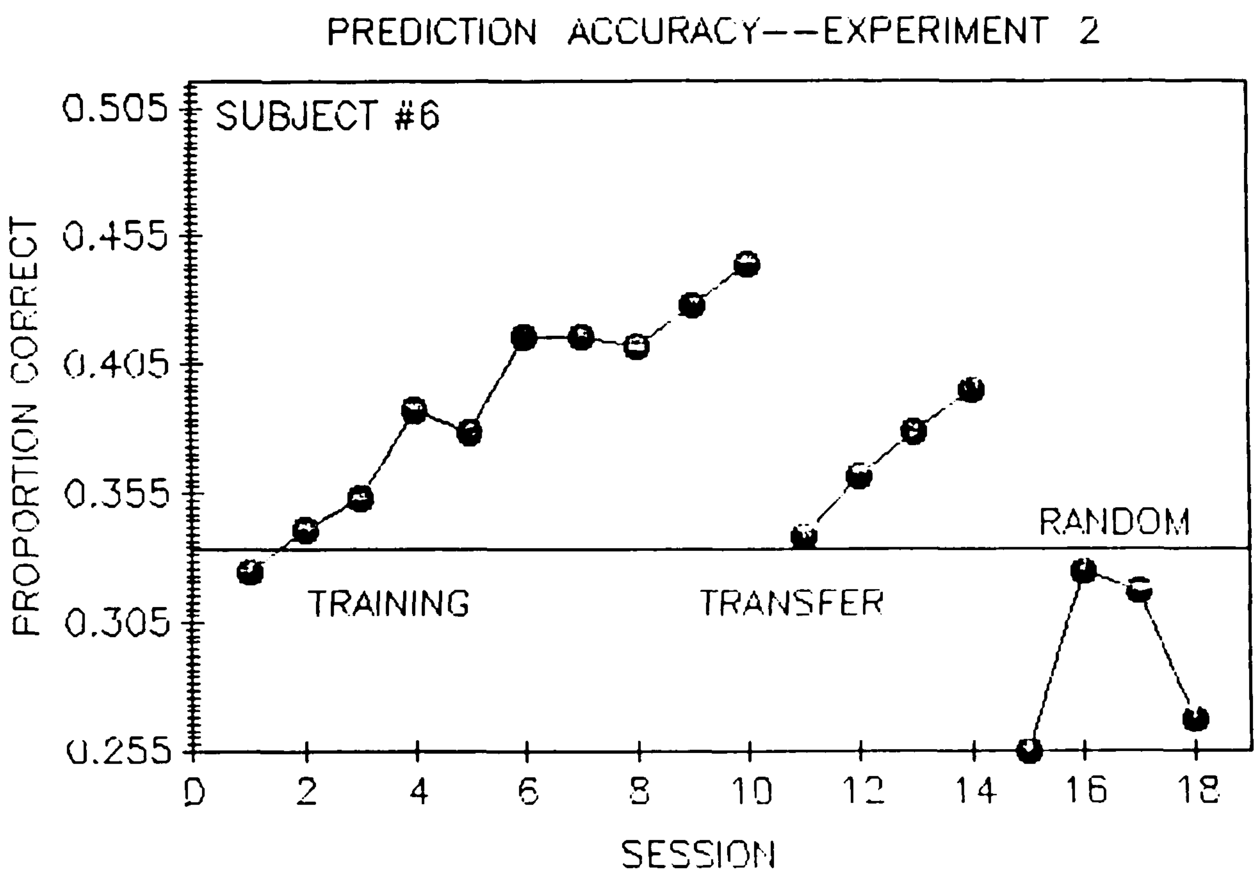

Figure 15: Prediction accuracy, subject 6, Experiment 2.

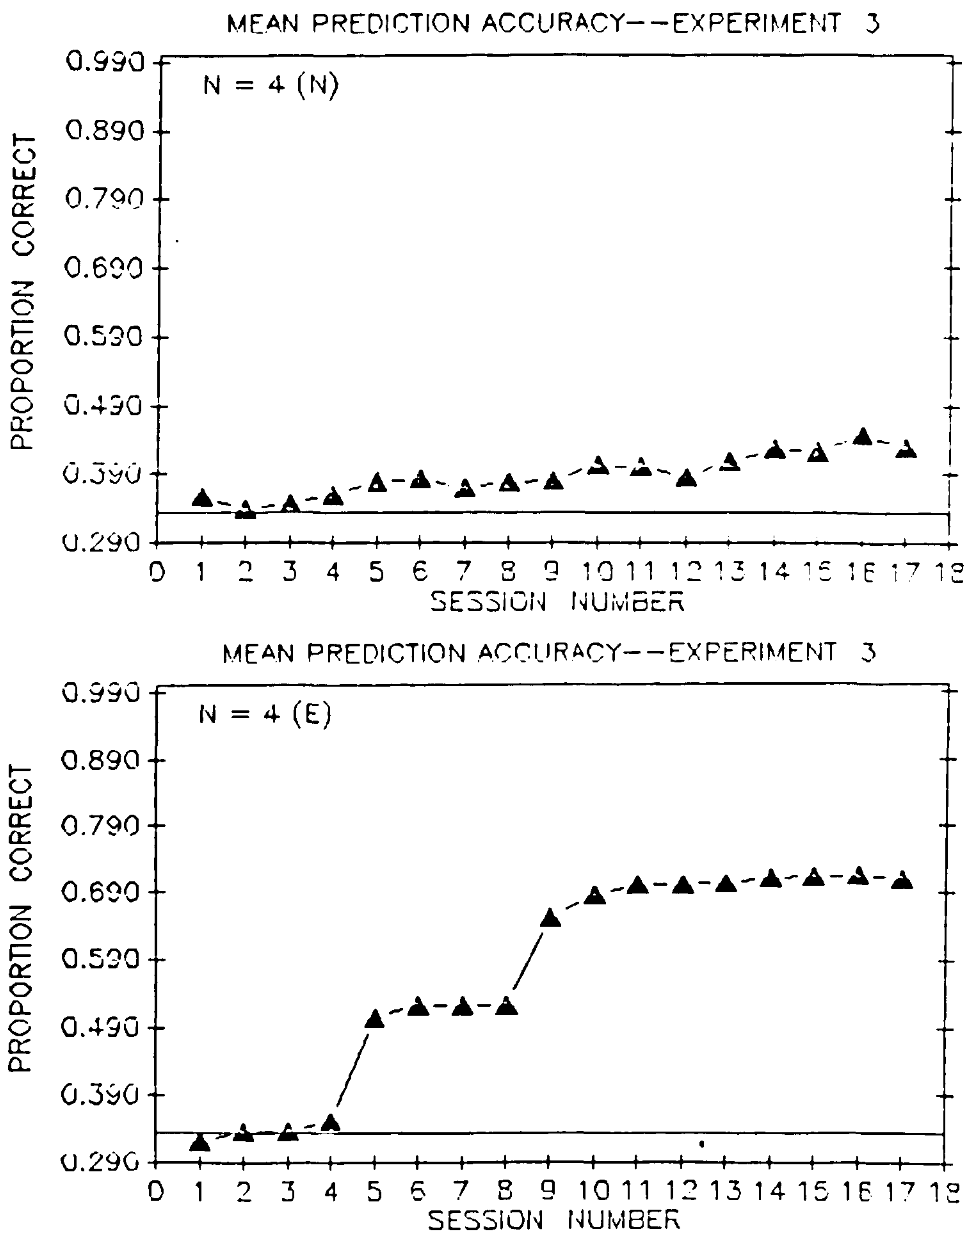

Figure 16: Average prediction accuracy, Experiment 3.

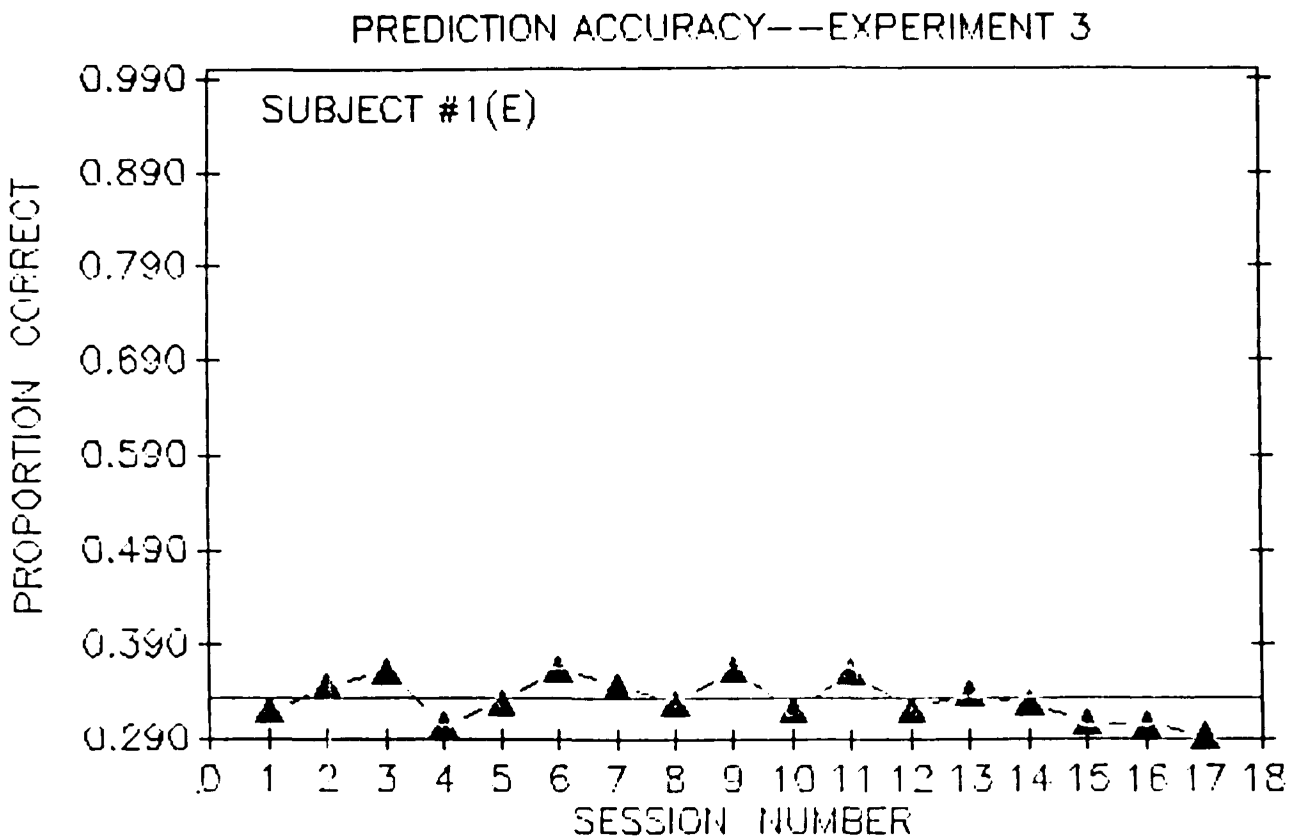

Figure 17: Prediction accuracy, subject 1(E), Experiment 3.

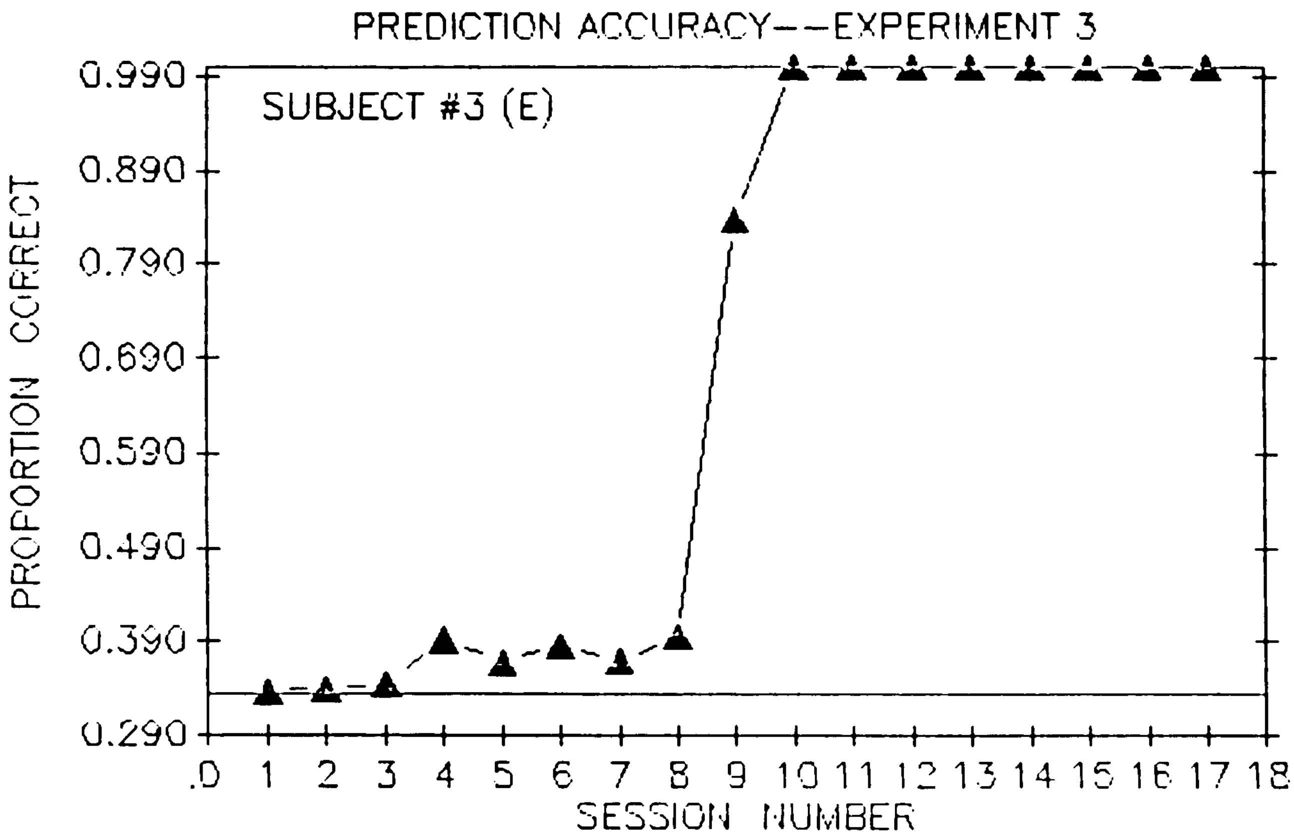

Figure 18: Prediction accuracy, subject 3(E), Experiment 3.

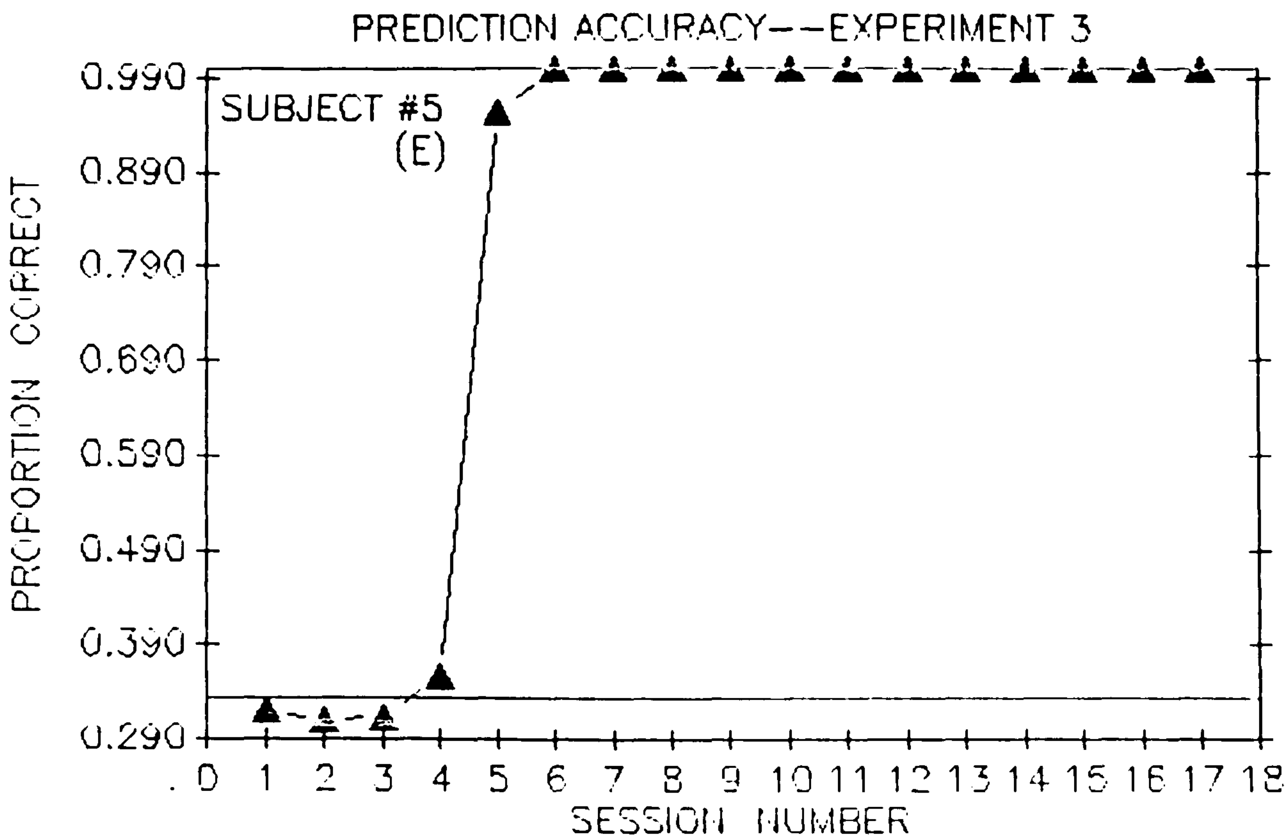

Figure 19: Prediction accuracy, subject 5(E), Experiment 3.

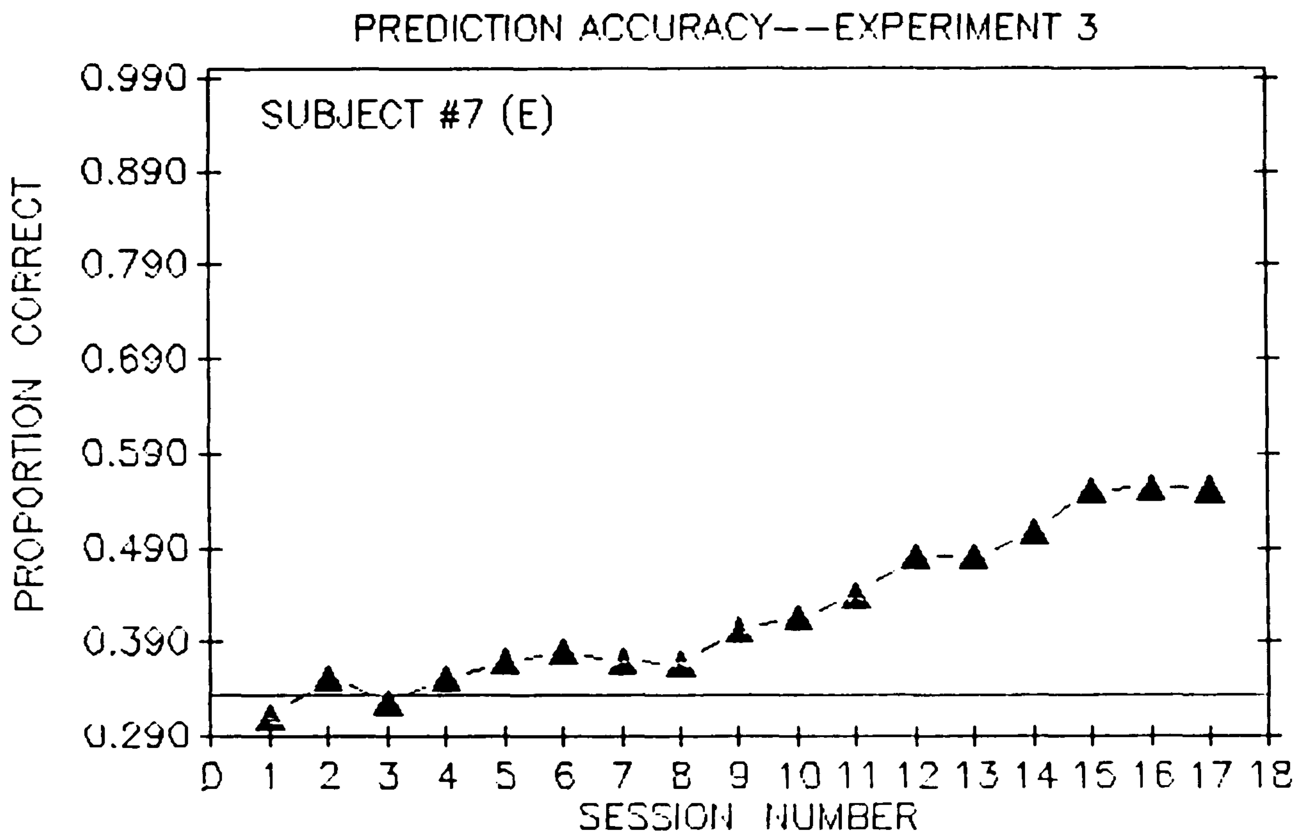

Figure 20: Prediction accuracy, subject 7(E), Experiment 3.

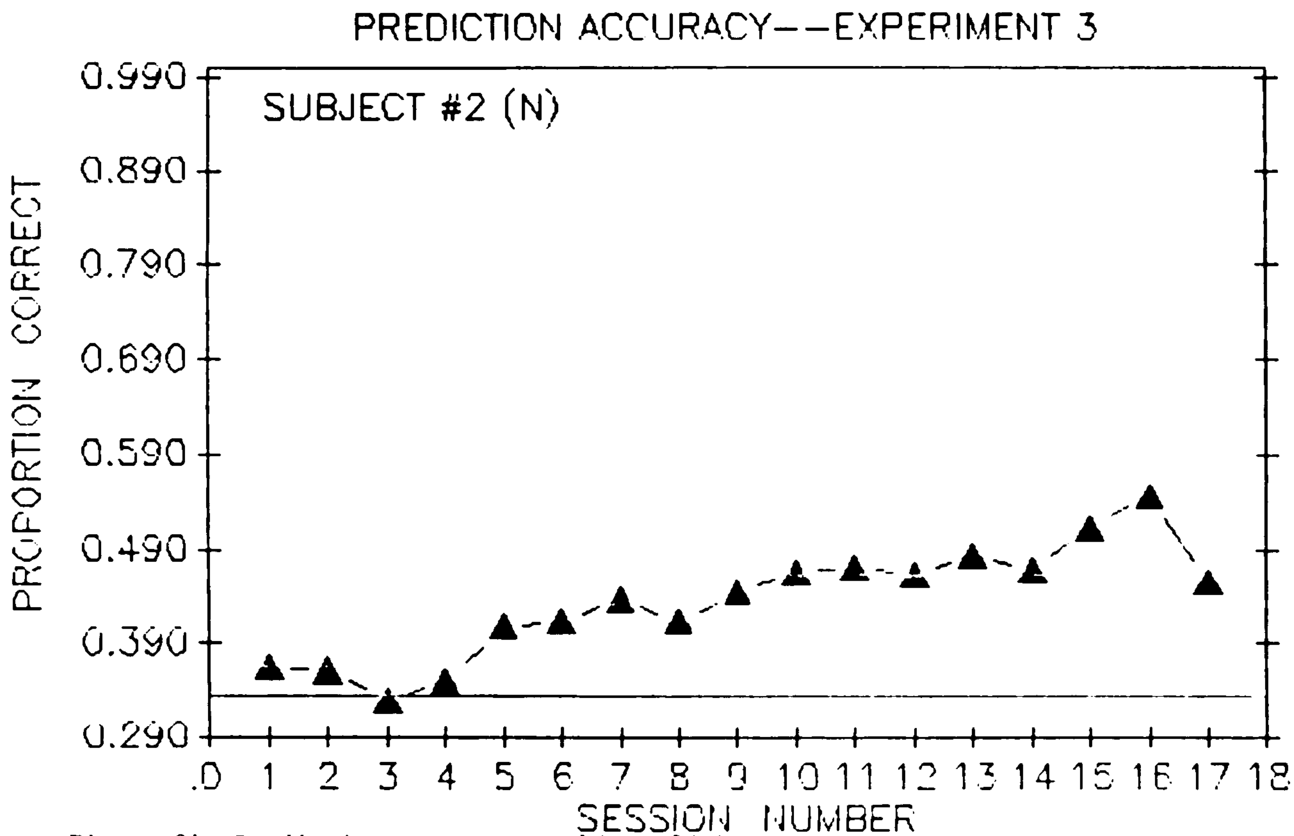

Figure 21: Prediction accuracy, subject 2(N), Experiment 3.

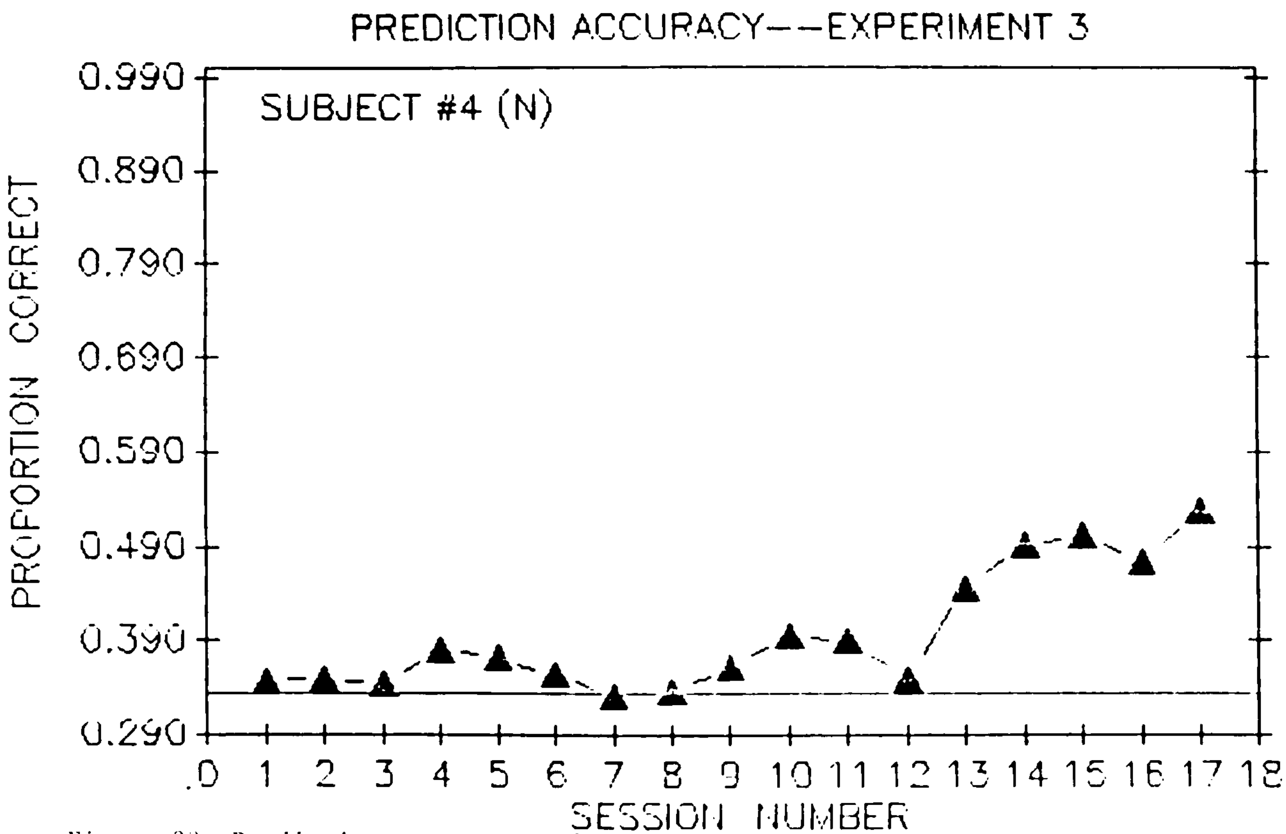

Figure 22: Prediction accuracy, subject 4(N), Experiment 3.

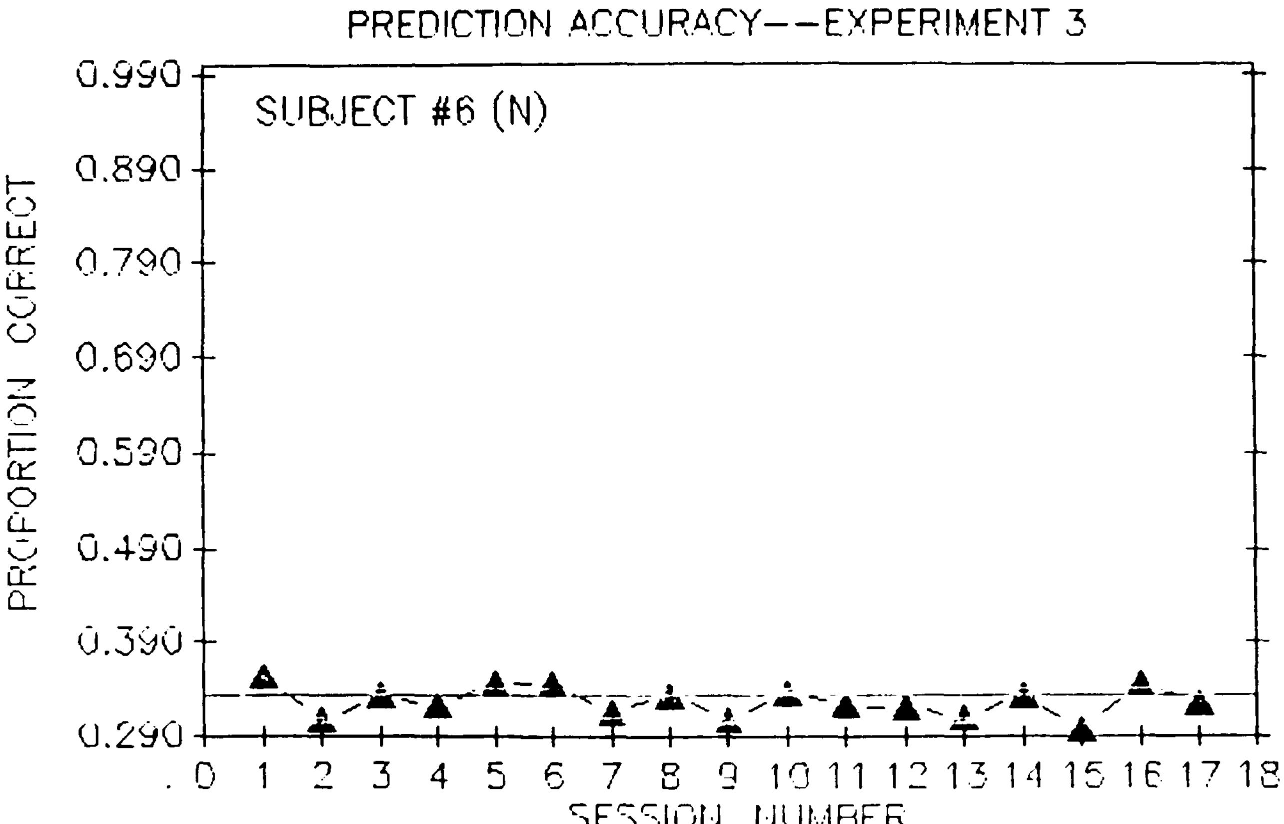

Figure 23: Prediction accuracy, subject 6(N), Experiment 3.

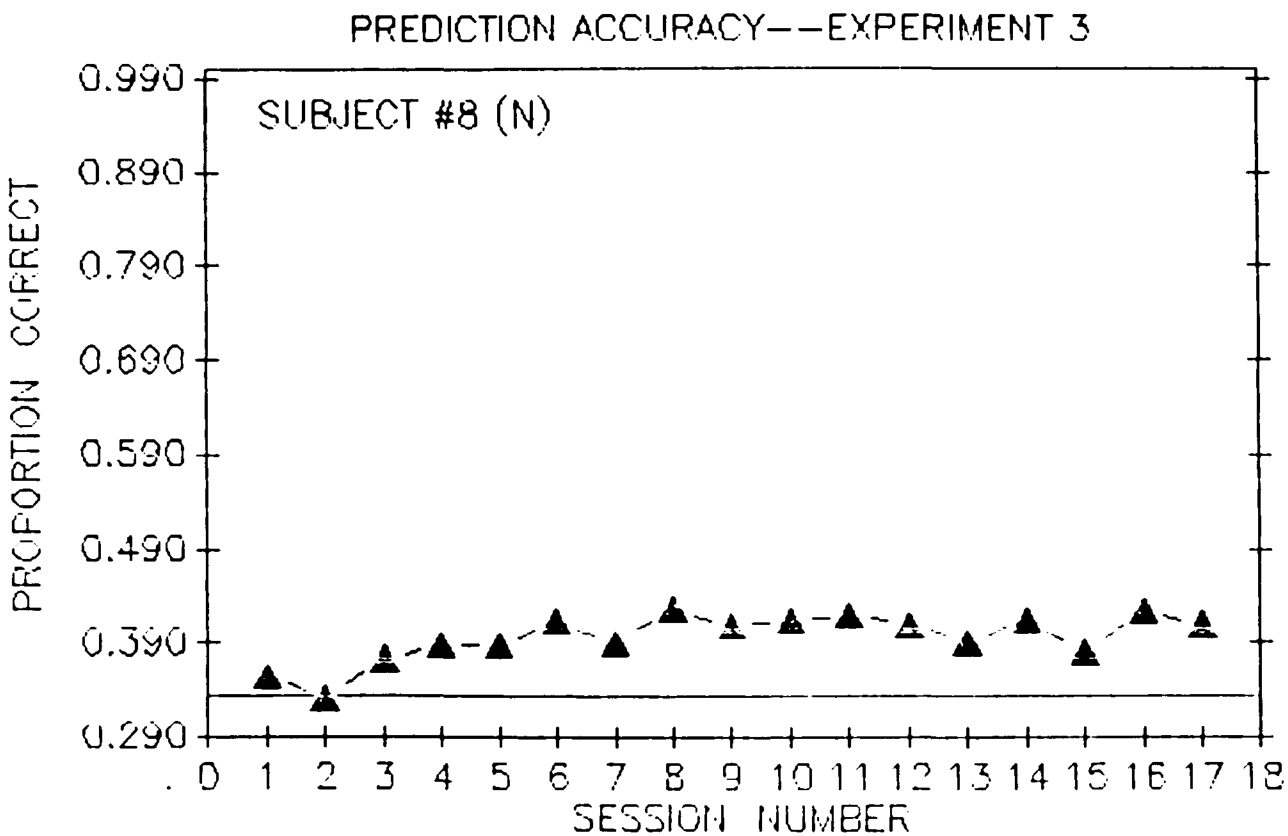

Figure 24: Prediction accuracy, subject 8(N), Experiment 3.

APPENDIX

In order to explore more deeply the possibilities that subjects are "picking up on" salient patterns, and that accuracy in predicting those patterns accounts for a significant portion of subjects' overall accuracy, two additional analyses are presented in this Appendix.

The first analysis is an extension of those analyses that attempted to explore <u>which</u> patterns could potentially be highly salient. Recall that the continuous patterns and the single-alternating patterns were found to fit that description in that they were predicted correctly on a considerable proportion of trials (see Table 10); they were also spontaneously mentioned by some subjects as being very memorable. However, there is a possibility that those sequences that have four spots in any one box (the "quads") might also be highly salient. There are 33 such sequences; three of them are the continuous patterns, previously explored, which have <u>five</u> of the same locations in a row. The remaining 30 patterns are listed in Table A-1. Note that 18 of the 30 predict Box 1 as a correct response, 6 of the 30 predict Box 2, and the remaining 6 predict Box 3.

Table A-2 shows the resulting proportion of correct responses for each subject in Experiment 2, Phase I, and in Experiment 3, when these 30 additional sequences are removed from the analysis of accuracy (that is, the figures report the accuracy without the 30 "quads" and without the 9 highly salient patterns noted in Chapter 3). All of these figures are significantly above the proportion of correct responses that would be expected by chance alone in the ensuing reduced number of trials. The normal approximation to the binomial distribution indicates that in the hypothetical shorter experiments, a subject would have to achieve at least 35.1% correct responses for Experiment 2, Phase I, and at least 34.1% correct responses for Experiment 3, in order to have predicted at a significantly above-chance level.

The second analysis considers the pattern of responses emitted by the subjects considered as a function of the correct responses. Table A-3 shows these data for Experiment 2, Phase I, and for Experiment 3, along with the probability of each response being emitted and the proportion of total correct responses accounted for by correct responses to each of the three types of sequences. Once again, this analysis reveals no inordinate "skew" in subjects' responses that would suggest a simplified prediction strategy.

REFERENCES

Berry, D.C., & Broadbent, D.E. (1984). On the relationship between task performance and associated verbalizable knowledge. Quarterly Journal of Experimental Psychology, 36, 209-231.

Brooks, L.R. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch and B.B. Lloyd (Eds.), Cognition and categorization, p. 169-211. New York: Wiley.

Cleeremans, A., & McClelland, J.L. (1991). Learning the structure of event sequences. Journal of Experimental Psychology (in press).

Cleeremans, A., Servan-Schreiber, D., & McClelland, J.L. (1990). Graded state machines: The representation of temporal contingencies in simple recurrent networks. Machine Learning, 2, 98-109.

Dulany, D.E., Carlson, R.A., & Dewey, G.I. (1984). A case of syntactical learning and judgement: How conscious and how abstract? Journal of Experimental Psychology: General, 113, 541-555.

Dulany, D.E., Carlson, R.A., & Dewey, G.I. (1985). On consciousness in syntactical learning and judgement: A reply to Reber, Allen, and Regan. Journal of Experimental Psychology: General, 114, 25-32.

Hasher, L., & Chromiak, W. (1977). The processing of frequency information: An automatic mechanism? Journal of Verbal Learning and Verbal Behavior, 16, 173-184.

Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. Journal of Verbal Learning and Verbal Behavior, 16, 107-112.

Hasher, L., & Zacks, R.T. (1979). Automatic and effortful processes in memory. Journal of Experimental Psychology: General, 108, 356-388.

Hasher, L., & Zacks, R.T. (1984). Automatic processing of fundamental information. American Psychologist, 39, 1372-1388.

Kushner, M., Cleeremans, A., & Reber, A. (1991). Implicit detection of event interdependencies and a PDP model of the process. In Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society, 215-220. Hillsdale, NJ: Lawrence Erlbaum.

Lewicki, P. (1986). Nonconscious social information processing. New York: Academic Press.

Lewicki, P. Czyzewska, M., & Hoffman, H. (1987). Processing information about covariations that cannot be articulated. Journal of Experimental Psychology: Learning, Memory, and Cognition, 12, 135-146.

Lewicki, P., & Hill, T. (1989). On the status of nonconscious processes in human cognition: Comment on Reber. Journal of Experimental Psychology: General, 118, 239-241.

Lewicki, P., Hill, T., & Bizot, E. (1988). Acquisition of procedural knowledge about a pattern of stimuli that cannot be articulated. Cognitive Psychology, 20, 24-37.

Mathews, R.C., Buss, R.R., Stanley, W.B., Blanchard-Fields, F., Cho, J.-R., & Druhan, B. (1989). The role of implicit and explicit processes in learning from examples: A synergistic effect. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 1083-1100.

Millward, R.B., & Reber, A.S. (1972). Probability learning: Contingent event sequences with lags. American Journal of Psychology, 85, 81-98.

Nissen, M.J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. Cognitive Psychology, 19, 1-32.

Perruchet, P., Gallego, G., & Savy, I. (1990). A critical reappraisal of the evidence for unconscious abstraction of deterministic rules in complex experimental situations. Cognitive Psychology, 22, 493-516.

Reber, A.S. (1967). Implicit learning of artificial grammars. Journal of Verbal Learning and Verbal Behavior, 6, 855-863.

Reber, A.S. (1969). Transfer of syntactic structure in synthetic languages. Journal of Experimental Psychology, 81, 115-119.

Reber, A.S. (1976). Implicit learning of synthetic languages: The role of instructional set. Journal of Experimental Psychology: Human Learning and Memory, 2, 88-94.

Reber, A.S. (1989). Implicit learning and tacit knowledge. Journal of Experimental Psychology: General, 118, 219-235.

Reber, A.S., & Millward, R.B. (1968). Event observation in probability learning. Journal of Experimental Psychology, 77, 317-327.

Reber, A.S., & Millward, R.B. (1971). Event tracking in probability learning. American Journal of Psychology, 84, 85-99.

Reber, A.S., Walkenfeld, F.F., & Hernstadt, R. (1991). Implicit and explicit learning: Individual differences and IQ. Journal of Experimental Psychology: Learning, Memory, and Cognition, 17, 888-896.

Scarborough, D.L. (1989). Personal communication.

Simon, H.A., & Sumner, R.K. (1968). Patterns in music. In B. Kleinmutz (Ed.), Formal representation of human judgement. New York: Wiley.

Willingham, D.B., Nissen, M.J., & Bullemer, P. (1989). On the development of procedural knowledge. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 1047-1060.

Zacks, R.T., Hasher, L., & Sanft, H. (1982). Automatic encoding of event frequency: Further findings. Journal of Experimental Psychology: Learning, Memory, and Cognition, 8, 102-116.

Zechmeister, E., & Nyberg, S. (1982). Human Memory. Monterey, CA: Brooks-Cole.