

Hobbits and Orcs: Acquisition of a Sequential Concept¹

JAMES G. GREENO
The University of Michigan

Subjects solved the hobbits-orcs problem repeatedly until they made no errors on two successive solutions. Groups varied in the kind of feedback given after errors. Further support was obtained for Thomas' (this journal) conclusion that subjects organize the sequence of moves in small subsequences, rather than individual moves. Differences between feedback groups suggested that the organization is oriented forward in the problem, and this contrasts with GPS' retroactive organization of this problem. Comparison between feedback groups also leads to the conclusion that subjects learned from positive information indicating which response is correct, rather than by a process of eliminating errors or sampling new strategies after errors. Analysis of acquisition data agreed with the hypothesis of all-or-none learning at individual states, except for one state where the additional complexity was related to an ambiguity about backward moves. The general pattern of results was invariant over a change in the characters of the problem (men and elves) designed to produce a reversal of relationship between the characters concerning who was prevented from outnumbering whom.

Three distinguishable but interrelated theories are involved in sequential problems and concepts. One is a theory of problem solving, dealing with the process of finding a correct sequence of actions when a problem is presented. If a person has had extensive experience in a situation, a correct sequence of actions may be well learned. A second theory describes performance after a sequential concept has been acquired, specifying the structure of the acquired concept, and processes of memory retrieval and application of the person's knowledge. And a third theoretical problem is to analyze the process of transition through which a person acquires a sequential concept, learning how to get something done without needing to solve a problem.

The purpose of this paper is to present results that mainly bear on the third theoretical question—how is a sequential concept acquired through experience in a problem-solving situation? The task studied was the hobbits-orcs problem developed as an experimental task by Thomas (1971).

¹ This research was sponsored by the National Science Foundation under grant GB31045. I am grateful to John C. Thomas, Jr., for several suggestions regarding this report. Requests for reprints should be sent to James G. Greeno, University of Michigan, Ann Arbor, Mich. 48104.

While Thomas studied initial solution of the problem and part-to-whole transfer, subjects in the present experiments carried out a solution of the problem, then were returned to the beginning state 331 and solved the problem again, repeating the process until a criterion of two successive errorless solutions was achieved. Observation of performance during these learning trials should be informative regarding the process of learning the sequence of moves involved in solving the problem. In addition, these experiments provided additional data to those of Thomas regarding the places in the problem where subjects have difficulty and, to some extent, why.

In order to provide information about the relative effectiveness of trial and error as compared to informative feedback, subjects in the first experiment were run in three groups, varying in the kind of feedback given after errors. One group, Noncorrection Allow Backward, were given freedom to retrace parts of the problem, and were informed only after errors that allowed orcs to eat hobbits. A second group, Noncorrection Prevent Backward, were informed after errors allowing eating and also after moves that would produce backtracking. These subjects had to produce the correct response before proceeding. A third group, Correction, were informed of both these kinds of errors, but then were told which response was correct rather than being told to try again.

The second experiment was run to check a possible source of difficulty in an early state of the problem, and to replicate the main findings of the first experiment using a different response mode and a changed display. The main change in the second experiment was a semantic one—the problem was described using men and elves as characters, rather than hobbits and orcs. There is a difficulty in the problem at state 321 (see Thomas' Fig. 2 in this journal) where the subject has to move two orcs. This results in having all the orcs together on the far side of the river with the boat. Since the orcs are characterized as "bad guys" in the problem, a source of difficulty might be reluctance to trust the orcs to bring the boat back and permit any hobbits to cross. (In Thomas' (1971) postexperimental questionnaire, some subjects remarked that the problem became easier after they learned they could "trust the orcs.")

To check this possibility, a single group of subjects was run with the cover story that three men and three elves had to cross the river. The men were well intentioned, but the elves were nervous creatures so that if any elves found themselves outnumbered by men they would cause themselves to disappear magically. The intent of the change was to remove the potential lack of trust in the creatures who must all be taken across the river first. In this version, the first three moves have to put all three men across the river with the boat, but it seems much less likely that

subjects would be reluctant to do this out of a lack of trust for the men. In addition to the change in the cover story, there were some procedural changes involving the mode of response and display. These will be described below.

METHOD

Subjects were volunteers from the Human Performance Center subject pool who were paid for their participation. Forty-nine, 50, and 51 subjects were run in the three groups of the first experiment.

The procedure for the Noncorrection Allow Backward group was virtually identical to the procedure used by Thomas (this journal) in his Control groups. Recall that this included feedback following moves that would let orcs eat hobbits—"You let orcs eat hobbits. Try again," after which the subject's response was removed and the subject had to give another response. After a backward move the situation was changed in accordance with the subject's response.

In the Noncorrection Prevent Backward condition, subjects were given the message about moves that let orcs eat hobbits. In addition, if the subject made a response that would be a backward move he saw, "That would take you backward. Try again." The subject's response was removed and the subject was to respond again.

In the Correction group, "Try again" was not part of the feedback message given after errors. After errors that would allow eating or making a backward move, the message "You let orcs eat hobbits," or "That would take you backward" was followed by the word "Move," along with the correct response. For example, if the subject was in state 321 and gave the response 1H, the subject saw, "That would take you backward. Move 20." This feedback message remained on the screen for about 4 sec, then the display was changed to the situation that the correct move would produce. If an error was made in state 331, the indicated correct move was always two orcs.

As in Thomas' study, subjects in the Noncorrection Allow Backward group had the option of restarting the problem by entering the word RESTART instead of a move. This option was not available in the other condition.

As a means of pacing subjects' responses, mainly to avoid excessive long-range planning, the message "Hurry please" appeared on the screen if the subject did not respond for 15 sec after a display appeared. Instructions to subjects explained that this message was only a reminder, and fast responding was not required to meet the criterion. Subjects typically saw the "Hurry" message for the first two or three responses, and rarely thereafter.

In the second experiment, a single group of 66 subjects was run using the cover story about men and elves, and the Noncorrection Prevent Backward procedure. Instead of typing the desired move on the keyboard and entering it, subjects responded using a row of four buttons marked M, E, Row, and Oops. The display had M's and E's denoting the men and elves on the two sides of the river, but instead of the word BOAT, the symbol $< >$ appeared on the side where the boat was. When the subject pushed the button marked M, one of the M's was taken away from the set appearing on the side with the boat, and an M appeared inside the "boat." Thus, pushing the button marked M had the effect of putting one of the men into the boat. Similarly, the subject could put an elf into the boat by pushing the button marked E. When the subject had the boat loaded as he wished, he pushed the button marked Row. This caused the move to be evaluated by the computer system. If the move was legal and not backward, the boat and its contents moved across the screen (actually, jumped across in three leaps) and the men and elves in the boat got out on the other side. If the move was not legal or would be backward, an appropriate message was given to the subject, the men and elves in the boat returned to the bank, and the subject was told to try another move. Any time that the subject felt he had put someone in the boat wrongly, he could push the button marked Oops, which returned all the M's and E's currently in the boat to the shore, allowing the subject to restart that move.

The major effect of this change in procedure was to provide more information visually about the effect of a move before the subject actually entered the move. In the first experiment, as in Thomas' study, the subject typed in the numbers and letters indicating what was to be moved, but the display was not changed as this occurred. This meant that the subject did not see the effect of removing creatures from the bank where the boat was. In the procedure of the second experiment, the creatures that would be left after the move could be seen by the subject prior to actually entering the move by pushing the Row button.

An additional procedural change was to keep subjects working for a longer criterion of correct solutions. Subjects continued until they had given a total of 12 errorless solutions.

RESULTS

In this report, only errors of the kinds that allowed orcs to eat hobbits (or elves to disappear) or that were backward moves are considered. Other errors, such as trying to move more hobbits or orcs than actually were on the side with the boat, neglecting to enter a number of creatures, etc. were considered to be of minor interest and were ignored.

Comparison of Feedback Conditions

Table 1 shows some relatively gross data regarding the speed of learning in the three conditions of the first experiment. Reliability was assessed using orthogonal comparisons: Noncorrection Allow Backward vs the other two groups and Noncorrection Prevent Backward vs Correction. The significant differences were those between Noncorrection Allow Backward and the other groups in the number of backward moves ($t(147) = 3.59, p < .002$), and the number of precriterion solutions ($t(147) = 3.62, p < .002$). These facts are explained by an hypothesis by Thomas (1971) that subjects move backward through the problem to give themselves more experience in situations that are puzzling. A second possible reason for more backward moves is that when backward moves were allowed subjects lacked feedback about backward moves and so may have been less clear that these moves took them away from the goal. Whether the considerable number of backward moves carried out by Noncorrection Allow Backward subjects were intentional or not, the additional experience acquired because of the retracing was beneficial, reducing the total number of times subjects had to solve the problem before learning the solution sequence.

The differences in errors between correction feedback and noncorrection feedback with backward moves prevented were not significant (backward: $t(99) = 1.13, .20 < p < .30$; eat: $t(99) = 1.07, .20 < p < .30$) so a firm conclusion is not warranted. If the numbers are taken at face value, they say that on the average, one time through the problem produced about as much learning for a Noncorrection Prevent Backward subject as for a Correction subject, and that the smaller number of errors made by Correction subjects was just due to the procedure that prevented them from making more than one error at any one state.

The results obtained in comparing the feedback conditions speak to a theoretical question about the learning process, sometimes called the question of the effective learning event (cf. Bower & Trabasso, 1964). The assumption is that whatever it is that causes learning caused it in all three conditions, and with about equal effectiveness. Differences be-

TABLE 1
Mean Errors and Number of Solutions before Criterion on the Problem

Condition	Total backward	Total eat	Solutions
Noncorrection, Allow Backward	12.7	9.6	3.6
Noncorrection, Prevent Backward	6.7	9.9	5.8
Correction	4.7	6.6	5.3

tween groups should be due to different frequencies of the effective learning event, brought about by the differences in procedure.

One hypothesis can be quickly rejected. The effective learning event did not correspond to a complete solution of the problem. If it did, then all groups should have required the same number of solutions before reaching criterion. Clearly, learning took place with respect to units smaller than the whole problem, as one would expect.

Two more reasonable hypotheses can be compared, although not decisively. According to some theories, the effective learning event occurs whenever the subject finds that an error has been made. One theory with this property is the idea that learning in problem solving is mainly a process of eliminating wrong responses or wrong strategies (Harlow, 1959). Another theory is that learning involves selection from an invariant set of strategies or hypotheses, and sampling occurs each time an error is detected (Restle, 1962).

An alternative hypothesis is that the effective learning event corresponds to obtaining information about what response is correct at a state where the subject does not yet know the correct response. One theory that leads to this hypothesis is that responses made before a subject knows the correct response are "best guesses," and the subject is informed by a successful outcome as much as by information that the response was wrong.

Of these two hypothesis, the one more favored by the data is that subjects learned from information about which responses are correct at the various states. The conclusion is not decisive, because it depends partly on the difference between number of errors made by Correction and Noncorrection Prevent Backward subjects and the difference was not significant. But, if each error gave an opportunity to learn, then the Noncorrection Prevent Backward subjects, making more errors during each solution of the problem, should have learned in a smaller number of solutions than the Correction subjects. The pattern of results obtained suggests that the additional errors made in the Noncorrection condition were not informative for the subjects. On the other hand, each time a Correction subject solved the problem, he received the same information about the correct responses as did a Noncorrection Prevent Backward subject. Subjects in both these conditions found out the correct response at each state just once each time they solved the problem. The close agreement between numbers of precriterion solutions is suggestive evidence that the effective learning event may have been the presentation of information about which response was correct at a state where the subject did not yet know the correct response.

It would be convenient to have a measure of difficulty that allows the

three conditions to be compared state by state as well as in overall performance. If the argument given here about the effective learning event is accepted, then neither the number of precriterion solutions nor the number of errors gives an unbiased comparison. The number of precriterion solutions is biased in favor of the Noncorrection Allow Backward condition because those subjects give themselves additional learning trials during a solution by retracing parts of the problem. The number of errors is biased against the Noncorrection Prevent Backward condition compared with the Correction condition, because the Noncorrection Prevent Backward subjects had to give the correct response to find out what it was, and this procedure led to their making more errors. The number of errors is even more strongly biased against the Noncorrection Allow Backward condition, because the retracing that was allowed had to be carried out using backward moves, which were counted as errors.

If the effective learning event occurred when a subject obtained information about a correct response, then unbiased comparisons can be made by counting a learning trial at a state each time the subject was informed of the correct response at that state. In the Correction condition, there was a single response at each state, and it was either correct or an error. In the Noncorrection Prevent Backward condition a single trial is counted for each state in each solution of the problem; the subject is correct if his first response was correct and an error is counted for the trial if any number of wrong responses occurred before the correct response was given. In the Noncorrection Allow Backward condition a trial is counted for a state each time the subject entered the state moving forward in the problem. For example, each time state 311 was entered from state 300 a trial was counted. If the subject's first response after such a forward entry was wrong, an error was counted. This scoring ignores errors that occurred when the subject entered a state moving backward, approximately to avoid counting as errors responses that the subject made in order to retrace part of the problem.

Results obtained with the above-described scoring method are in

TABLE 2
Mean Errors and Trials Before Criterion on Individual States, Summed Across States,
Counting a Trial at a State as All Responses up to a Correct Response at that
State (Responses After Back-ins Omitted)

Condition	Errors	Trials
Noncorrection, Allow Backward	8.3	13.7
Noncorrection, Prevent Backward	8.5	13.4
Correction	11.0	16.1

Table 2. To obtain these results, responses for each state were counted until a subject was correct in four consecutive trials. The data tabled here are the sums of the results obtained for the individual states.

Note that by defining trials in the way specified, scores for the two Noncorrection groups are nearly equal, and learning in the Correction group appears to have been a bit slower. In the next section, the apparent slower learning of the Correction group will be discussed in relation to the particular states where it occurred. The fact that the definition of a trial as occurrence or presentation of correct response gives nearly equal overall learning scores suggests that this definition is an appropriate one for analyzing state by state performance, and since this definition also has a plausible theoretical basis it will be used for the analyses to be presented subsequently.

Comparison of Difficulty at Different States

Figure 1 shows the mean number of errors in each state for the three conditions of the first experiment. The function graphed is the serial position curve of difficulty in learning the sequence of responses needed to solve the problem. The graph reveals that special difficulty occurred at states 321 and 110, in agreement with Thomas' conclusions about difficulty during initial problem solving (this journal, Fig. 2). Note again, that the hard states have the maximum number of alternative possible responses (five) but if anything the relation between number of alternatives and difficulty is even weaker in these learning data than in Thomas' observations of problem solving.

The patterns obtained for the three conditions seem different in two interesting ways. First, the Noncorrection Allow Backward subjects made more errors at state 321 and fewer errors at state 110 than the subjects

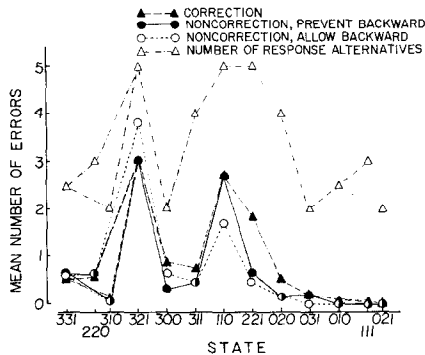


FIG. 1. Mean number of errors at each state in the first experiment.

in the other two conditions. This was probably due to retracing. Subjects who retraced the first part of the problem a few times had a few learning trials at state 321 before their first trial at state 110. This would have two effects, relative to the situation in the other conditions where a single trial occurred for both states 321 and 110 during each solution. First, if learning at state 321 produced positive transfer at state 110, there would be more positive transfer at 110 in the Noncorrection Allow Backward condition than in the other conditions. Conversely, if learning at state 110 produced positive transfer at state 321, there would be less positive transfer at 321 in Noncorrection Allow Backward than in the other conditions. The reason for this second effect is that trials at 321 pile up due to retracing and these added trials at 321 occur without possibility of benefit from learning at 110.

A second factor that might make state 321 harder for Noncorrection Allow Backward subjects is involved in the structure of the problem. State 321 is the single place in the problem where a backward move can be made without returning to the state just left on the previous move. Some evidence involving the effect of this structural feature on learning at state 321 will be presented later, under "Analysis of Acquisition." But this structural ambiguity could have had a special effect on the Noncorrection Allow Backward subjects, since they were not informed by feedback when they made backward moves. Analysis of the kinds of errors made at state 321 supported this conjecture. Counting only the first response made on each trial at 321, there were a total of 97, 120, and 99 errors allowing orcs to eat hobbits in the Correction, Noncorrection Prevent Backward, and Noncorrection Allow Backward conditions, respectively. The corresponding numbers for backward moves were 36, 49, and 87.

The second noteworthy fact about Fig. 1 is that while performance in the Correction condition was virtually identical to Noncorrection Prevent Backward in the two hardest states, it was noticeably worse in the relatively easy states that follow those states—especially the states following 110. Recall Table 2, where the Correction condition showed slightly worse performance than the other conditions. Figure 1 shows the decrement to have been located at relatively easy states in the problem, but at states that follow the states that are hard.

An explanation of this interaction is that subjects in the Noncorrection conditions engaged in some look-ahead at the hard states 321 and 110, and that the correction feedback disrupted this. Recall that in the Noncorrection conditions subjects had to continue responding until giving the correct response. In the Correction condition the correct response was given to the subject after an error occurred. Look-ahead at state 321

could take the form of a response like, "I can take two orcs across and bring one back." At state 110 a response might be, "I can take an hobbit and an orc across and come back with both hobbits." If a subject in Non-correction found such a response his next move would be easy. But subjects in Correction were not given as much opportunity to find these two- or three-move responses, being interrupted by the presentation of the correction feedback.

Recall that the second experiment was run with a cover story about men and elves to check the possibility that state 321 is difficult because subjects are reluctant to trust all the orcs with the boat. Figure 2 shows the serial position curve obtained in the second experiment. Apparently neither the change in semantic content of the problem nor the procedural change in the manner of entering responses made any substantial difference in the relative difficulty of the various states.

Another measure of relative difficulty in the different states is the time to respond. Figure 3 shows the mean time to respond in each state at different stages of learning in the three conditions of the first experiment. Panel (a) is for the first solution of the problem. Panel (e) has the data from each subject's two criterion trials. And panels (b), (c), and (d) have Vincent thirds of the solutions between the first solution and the criterion solutions. Subjects who had fewer than four solutions including the criterion were excluded from this analysis. The numbers of subjects included were: Noncorrection Allow Backward—41, Noncorrection Prevent Backward—44, Correction—42. For subjects with fewer than three solutions between the first solution and criterion, their second solution was entered two or three times as needed to provide data in all three Vincent thirds.

The main result in these reaction times is that time appears to be longest at three states: the first state, and the two states where subjects




FIG. 2. Mean number of errors at each state in the second experiment.

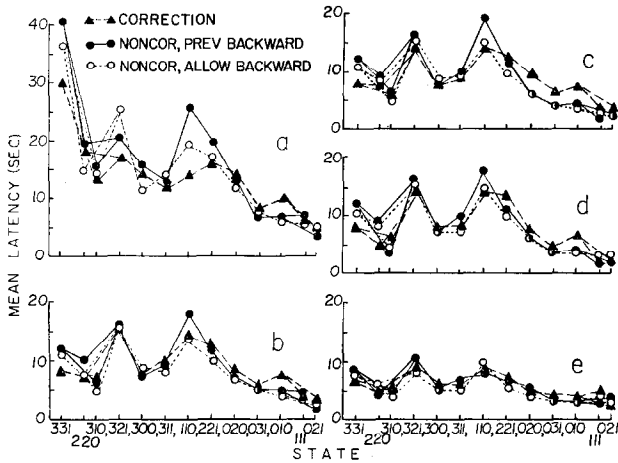


FIG. 3. Mean time to respond at each state at different stages of learning in the first experiment: a—first trial; b, c, d—Vincent thirds between first trial and criterion; e—criterion trials.

had difficulty learning the correct response. This three-peaked pattern was very clear even in the first solution for the Noncorrection Allow Backward condition. The pattern appeared almost as quickly for Noncorrection Prevention Backward. The pattern of reaction time differences was much less clear in the Correction condition, not appearing at all during the first solution and showing elevations at states 221 and 010 until the criterion trials. The relative lack of clarity in the reaction time pattern from the Correction condition provides further evidence that the correction feedback probably disrupted a tendency for subjects to respond in an organized way involving units of a few moves.

Figure 4 shows reaction time data from the second experiment. Reaction times for precriterion solutions are omitted. They were like those of the Noncorrection Prevent Backward group of the first experiment, except that differences between states were more pronounced in the second experiment. The three curves shown are for the mean time before pressing the Row button during subjects' first two errorless solutions (criterion), the subjects' last two errorless solutions (final), and solutions between these (between). There is a strong suggestion in these data that the reaction time function is flattening toward a non-descript asymptote. It is possible, of course, that the loss of pattern is due to over-all speeding and a consequent floor effect. But it is also reasonable to conclude that the reaction time differences observed during initial solution do not reflect structural features of the sequential concept that subjects have acquired after learning is complete.

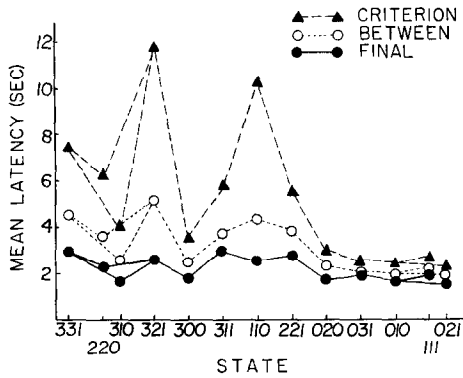


FIG. 4. Mean time to respond at each state after learning in the second experiment. Criterion trials correspond to panel e in Fig. 3.

Comparison with GPS—The data showing the relative difficulty at different states allows a comparison of human performance with performance by one version of GPS. Ernst and Newell (1969) presented a record of GPS' performance on the present problem. I have taken the number of operations carried out by GPS for each move as a measure of the difficulty GPS has at each state. This is plotted in Fig. 5, along with some data from the first experiment. I have chosen data from the Noncorrection Prevent Backward condition for this comparison because it seems to me to most closely approximate the constraints that are given to GPS for this problem. GPS never returns to an earlier goal, so the retracing shown by the Noncorrection Allow Backward condition never appears. On the other hand, GPS always produces each correct response, so the Correction condition seems less appropriate than a Noncorrection condition where subjects acquired the experience of developing a correct goal at each state.

One set of human data presented are the proportions of errors at each state in the initial solution. The other data are the times to respond during the criterion trials. It seems useful to think of GPS' performance in relation to both initial problem solving and to performance by human subjects who have learned the solution sequence. Since GPS does not make errors on this task, it could be that human problem solvers acquire important features of GPS' problem space through experience with the problem. In any case, the exact choice of human data for comparison with GPS seems relatively unimportant, since all of the data from human performance were quite similar in overall pattern. The data shown in Fig. 5 were arbitrarily scaled to give a visual range of points about equal to that plotted for GPS' performance.

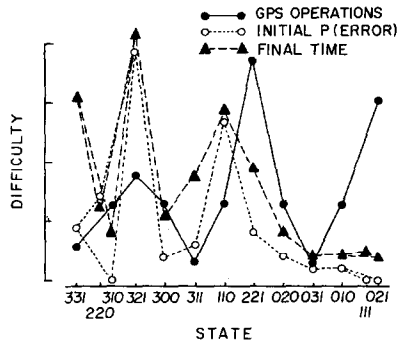


FIG. 5. Measures of difficulty at each state in Noncorrection Prevent Backward condition of the first experiment, compared with number of operations used by GPS at each state.

Figure 5 indicates that the places where GPS has difficulty and the places where human subjects have difficulty do not correspond very closely. GPS has its greatest difficulty at state 221, where humans take a moderate amount of time but make very few errors initially, and at state 021, which is trivially easy for human subjects. On the other hand, the hard states for humans, 321 and 110, are of only moderate difficulty for GPS.

A significant source of GPS' discrepancy from human performance is a result of the way in which its goal tree is constructed. When a goal is established, GPS evaluates the situation to see whether the goal can be accomplished. If it can not, there is an effort to modify the situation to permit accomplishment of the goal. If an operator is available that produces the needed modification, the operator is used and GPS then goes back to the previous goal. This procedure can create a difficulty that Quinlan and Hunt (1968) called the looping problem.

An example of how this system causes difficulty for GPS occurs at state 021, a state where 150 human subjects in the first experiment made a total of five errors. At state 010, GPS correctly recognizes that the problem is to transfer the last orc from the left side. The goal of moving one orc from the left to the right is established. However, an impediment is detected, since the boat is on the right side. This leads to establishing a subgoal of bringing the boat to the left, which can be accomplished by transferring an orc from the right to the left. GPS then reverts to the previous goal of transferring *one* orc from the left. Of course, human subjects immediately recognize that *two* orcs now need to be transferred, but GPS finds itself trying to undo what it has just done, and takes eight

operations by my count in finding out that it should try to take both of the remaining orcs across.

Analysis of Acquisition

The trial-by-trial data from each state in each condition were analyzed using a Markov model of learning. The model incorporates the hypothesis that learning the correct response at a state involves a simple, discrete all-or-none transition between states where the response is not known to a state where it is. For this application, three parameters were estimated: t , the probability of knowing the correct response at the beginning of the problem; c , the probability of learning the correct response on any trial; and p , the probability of giving the correct response on a trial when learning had not yet occurred. (See Atkinson, Bower & Crothers, 1965, or Restle & Greeno, 1970, for general discussion. Greeno & Scandura, 1966, gave an application where t was a free parameter, as in the present case.)

Trials were defined in the way specified earlier. The distribution of total errors and trial of last error was obtained for each state in each condition. These data were used to estimate parameters; sufficient statistics are the proportion of subjects giving no errors, and the conditional mean number of errors and trial of last error, given at least one error. Then the empirical distributions were compared with theoretical distributions derived from the all-or-none model.

The result of these analyses was that performance at every state except one was consistent with the all-or-none model. The test of all-or-none learning is not powerful unless learning is at least moderately difficult, so it is not surprising that data for the states with practically no errors agreed with predictions from the model. However, there were seven cases where sufficient errors occurred to permit chi square goodness-of-fit tests to the distributions of total errors and trial of last error, and where the predictions and data agreed. These were state 110 in all three conditions in the first experiment and in the second experiment, state 220 in the Noncorrection Prevent Backward condition and states 300 and 221 in the Correction condition. The goodness-of-fit chi squares statistics summed over all seven of these states were $\chi^2(17) = 17.09$ for total errors and $\chi^2(23) = 26.46$ for trial of last error. The theoretical and empirical distributions of total errors for state 110, averaged over the four experimental groups, are presented in Fig. 6.

The one state in which learning was apparently not all-or-none was state 321. In the first experiment goodness-of-fit chi square statistics approached significance in all three conditions and the sum was sig-




FIG. 6. Distribution of total errors at state 110 averaged over four experimental groups: all-or-none theory—connected dots; empirical—histogram.

nificant ($\chi^2(12) = 25.25, p < .025$). In the second experiment, the single experimental condition gave a significant discrepancy from the model ($\chi^2(3) = 11.21, p < .025$). The theoretical and empirical distributions deviated in similar ways in all four conditions. The average distributions are shown in Fig. 7, where it can be seen that the deviation is of the kind that is produced by a system more complex than all-or-none learning. Distributions of trial of last error in the first experiment were closer to predictions and the deviations were not significant. In the second experiment, this statistic deviated significantly from the model ($\chi^2(4) = 13.96, p < .01$). In any case, the result involving total errors is sufficient to reject the all-or-none model for state 321 in both experiments and conclude that at that state learning involved at least two stages.




FIG. 7. Distribution of total errors at state 321 averaged over four experimental groups: all-or-none theory—connected dots; empirical—histogram.

A number of possible reasons for the complexity of learning at state 321 were considered. One interesting possibility that was not supported in the data involves the fact that state 321 can be entered from either of two other states, 220 or 310. This raises the possibility that trials at state 321 were really a mixture of two kinds of trials, some in which the situation presented the stimulus: "321 after 220," and others involving the stimulus "321 after 310." The appropriate analysis in such a case would involve the so-called two-element model (Suppes & Ginsburg, 1963) in which learning is assumed to be all-or-none with respect to each stimulus element, but there are two stimulus elements so the learning observed in the whole situation is more complex. In most applications of this model, the stimulus elements are hypothetical and the mixture of trials is analyzed as a random process. In the present situation the two kinds of trials can be separated in the data. Accordingly, trials involving state 321 were separated into two sets—those in which the subject entered 321 from 220 and those in which the subject entered 321 from 310. If the two-element model were correct for this situation, each of these subsets of sequences should have properties consistent with the all-or-none model. This did not appear to be the case. Although statistically significant discrepancies were not obtained (probably due to losing power by subdividing already small sets of data), the distributions in each subset appeared to have about the same form as seen in Fig. 6 for all the data from state 321.

An alternative analysis did appear to reveal the nature of the complexity at state 321. Recall that at state 321, unlike all other states in the problem, it is possible to make a backward move without returning to the state just left on the previous move. This makes learning to avoid backward moves somewhat harder at state 321 than at other states, and it is possible that this difficulty could cause a second stage of learning to occur. Two items of data turned out to be consistent with this possibility. First, the relative proportion of errors that were backward moves were calculated for the first experiment, and were found to increase over trials. Counting all errors that occurred the proportions of backward moves were .42, .39, .44, .56, .56, and .58 on trials 1–6, and .44 after trial 6. Counting only the first response on each trial, the proportions of errors that were backward moves were .18, .34, and .39 for trials 1–3, .49 for trials 4–5, and .39 after trial 5. The tendency for relative frequency of backward moves to increase suggests that learning might have involved a two-stage process, with an intermediate state in which backward moves were predominant.

A second test is somewhat stronger. If learning at state 321 involved two stages with an intermediate state where only backward moves oc-




FIG. 8. Distribution of number of errors allowing orcs to eat hobbits in all groups of the first experiment: all-or-none theory—connected dots; empirical—histogram.

curred, then omitting backward moves from subjects' protocols should leave data that are consistent with the all-or-none learning model. This was indeed the case. For the analysis, trials in the various conditions were defined as before, but a correct response was recorded unless the subject made one or more moves that let orcs eat hobbits (or that let elves disappear). The mean numbers of trials with these errors were nearly identical in the three conditions of the first experiment (2.39, 2.10, 2.37) so these data were combined. The empirical and all-or-none theoretical distributions of total "eat" errors are shown in Fig. 8. The deviation was not statistically reliable ($\chi^2(5) = 8.39$, $10 < p < .20$). Results in the second experiment were similar, although the number




FIG. 9. Distribution of number of errors allowing elves to disappear in the second experiment: all-or-none theory—connected dots; empirical—histogram.

of errors allowing elves to disappear was somewhat larger, 3.32. The empirical distribution is compared with the all-or-none prediction in Fig. 9; the null hypothesis of agreement could not be rejected ($\chi^2(3) = 1.64$, $.30 < p < .40$).

DISCUSSION

The combination of results reported in Thomas' article (this journal) and those of the present study appear to provide some new understanding of processes involved in problem solving and in the acquisition of a sequential concept through experience in a problem-solving situation.

Organization of Problem Solving

Thomas' findings gave strong evidence that subject's organization of a task during problem solving is organized in units that extend beyond the states of a problem behavior graph. Additional evidence favoring this view was obtained in the learning study of the present investigation. The fact that subjects in the Correction condition of the first experiment had more difficulty in the relatively easy states following states 321 and 110 than the comparable Noncorrection conditions seems to indicate that subjects in the Noncorrection conditions were organizing their responses at these difficult states in a way that involved the next state or so in the sequence.

If this inference is correct and human subjects do organize their effort to solve problems in relation to two or more moves in the future, this represents an interesting discrepancy between human performance and the performance shown by GPS. GPS has superhuman memory for its past activity in trying to solve a problem. But its effort to make further progress is confined to an attempt to find an operation that will reduce the most important difference that it finds between the current situation and the goal. In this sense, GPS' forward progress is locked to the problem behavior graph, because it cannot carry out operations that develop new opportunities. GPS therefore has an organization of the task that uses the problem-solving history in a strong way, but uses look-ahead in a way that apparently is inferior to the problem-solving performance of human subjects.

The looping problem which is encountered by GPS in the hobbits-orcs task is a case in point. GPS gets to state 021 from state 010 by saying, in effect, "I want to move an orc across (now) but I can't, so I'll move the boat to the left side." Once the boat has been moved to the left, GPS returns to its earlier goal of moving one orc from the left to the right side, but this goal is no longer appropriate. Human subjects in state 010 apparently say something like, "If I move an orc across to the left, I can

bring both orcs back." In other words, at state 010, human subjects decide on a pair of moves that will complete the problem.

That GPS organizes the task in a way differently from human subjects should not be interpreted as weakening the contribution made by the development of GPS to the theory of human problem solving. The fundamental orientation of GPS as a problem solver oriented toward reducing differences between present and desired situations seems sure to be a basic feature of human problem solving, as is borne out by the analyses given by Newell and Simon (1972). Another feature of GPS that has made a substantial impact on psychological theory and that seems virtually sure to be valid is the idea of multilevel processing (cf. Reitman, 1965). GPS solves problems utilizing a process of executive control that establishes subgoals, and selects methods for reducing differences between present and desired situations. The operations by which methods are carried out take place at another, lower level of processing. There seems little doubt that something like this distinction between general executive procedures and operations that are more specific to individual tasks is a correct component of a theory of cognitive processing. On the other hand, there is an important distinction between the organization of the process and the organization of the task. The cognitive structure corresponding to GPS' organization of a task is essentially linear, focusing on properties of the present situation and differences between those and desired properties and utilizing only very weak look-ahead properties. The evidence of these experiments suggests that a more complex cognitive organization of tasks may be characteristic of human problem solving.

A question about problem-solving organization that is orthogonal to the comparison between human and GPS performance is the question of how important the semantic content of the task environment is in problem solving. The difficulty at state 321 found by Thomas and confirmed in the present study was unanticipated, and it seemed plausible that a strong contributing factor was subjects' interpretation of the story and additions of semantic meaning to the formal aspects of the problem. A relatively strong effort to reverse that interpretation by phrasing the problem using men and elves as characters did not remove the difficulty at 321. Therefore, it seems reasonable to conclude that the difficulty there is syntactic, involving the formal aspects of the problem, rather than the semantics. It may also be remarked that the rather substantial procedural change involved in the men-and-elves experiment made no apparent difference in the pattern of errors and reaction times in the various states. This indicates that the organization of the problem in Thomas' experiments and in the first experiment of this article was prob-

ably not influenced to any great extent by the procedure of typing descriptions of the moves on a keyboard.

Process of Acquisition

There are three issues that seem to be informed by the data obtained regarding the acquisition of a sequential concept during problem solving. One is the nature of psychological changes that occur during this kind of problem solving, and the evidence on this point mainly relates to the kind of intratask transfer that apparently occurs. A second issue is the question of the effective learning event, and the third issue is the question whether changes that occur during learning are discrete or continuous.

The first issue is the most subtle, and inferences regarding it are the most indirect. The main fact involved is the pattern of differences found in the first experiment reported in this paper between Noncorrection Allow Backward and the other groups particularly at states 321 and 110. Subjects who retraced the early part of the problem (the Noncorrection Allow Backward subjects) took more trials to learn the response at state 321 and fewer to learn at 110, and it seems reasonable to conclude that there was something in common between these states. Thomas' results contain half of this finding—the Part-Whole subjects made fewer errors at 321 than did the control subjects who started the problem at the standard initial state 331.

This pattern of transfer is intriguing because at the level of overt responses or anything that comes to mind introspectively about responding at those two states, they are quite different. At 321 the subject must move two orcs. At 110 the required move is an orc and an hobbit. At 321 the hobbits and orcs have to be put on separate sides of the river. At 110 the numbers of hobbits and orcs on the two sides have to be kept equal. One feature that is shared by these two states is a need to check the side from which the boat is moving to make sure no more orcs than hobbits are being left there. However, this seems an unlikely source of transfer, since if this had been a dominant factor in the problem the procedural change involved in the men-and-elves experiment should have reduced the difficulty of states 321 and 110, because in the latter experiment subjects saw changes in the side being left as they loaded the boat. It seems likely, then, that the change leading to intratask transfer involved some rather general features of the problem space, as Thomas (this journal) also concluded.

Whatever the exact basis of positive transfer between states 321 and 110, it seems very likely that positive transfer did occur. The additional evidence of positive transfer in the present data makes Thomas' (this

journal) finding about state 110 in his Part-Whole group even more remarkable. Recall the finding was that the Part-Whole group made *fewer* errors starting at state 110 than were made by the Control group between the time they reached state 110 and completed the problem. The present data indicate that Thomas' Control group probably benefited from some positive transfer at state 110 from having worked past state 321. Thus, the *disadvantage* of the Control group relative to the Part-Whole group at state 110 was apparently large enough to overcome some positive transfer and *still* produce a significant increase in number of moves needed to complete the problem. The inference based on Thomas' data that difficulty at 110 is considerably increased by a context effect of some sort gains strength in the present results.

The other inferences about acquisition from these data seem more straightforward. The data seem to favor an hypothesis that subjects learn in this task on the basis of information about which response is correct at each state, rather than about errors that they make. And the evidence based on the Markov analysis allows the pleasant hypothesis that in this sequential problem-solving task, components of learning occur in an all-or-none fashion. This latter finding is congruent with Bjork's (1968) analysis of acquisition of arithmetic sequences.

Structure of the Acquired Concept

Previous analyses of serial pattern learning have used error profiles (Restle & Brown, 1970) and patterns of reaction time (Pollio & Reinhart, 1970) to induce structural properties of the acquired sequential concept. The method does not seem to be valid in the present case, and the reason is theoretically interesting. First, the method identifies the entry points of structural units with difficult places in the problem. This would lead to grouping moves in the following units: (20, 10) (20, 10, 2H) (1H10, 2H, 10, 20, 10, 20). This is not a very natural way of organizing the problem, and it seems doubtful that subjects have this organization, especially in the light of the latencies shown by subjects in the men and elves experiment after many trials of overlearning. If the above-specified groupings were used, the problem would have the organization: (1) get an orc across and return, (2) get two hobbits and two orcs across, (3) get the rest across.² This seems a much less natural and useful organization of the problem than others, such as (1) get all the orcs across, (2)

² In earlier reports at the Mathematical Psychology meetings at Princeton, September, 1971 and at Indiana, April, 1972, I expressed my then-current opinion that this structure had actually been acquired. The latencies from the men-and-elves experiment and my inability to find a satisfying defense against objections on this point from anonymous reviewers have led me to an opposite view which is expressed.

get all the hobbits across, (3) get the orcs back across (cf. Amarel, 1968). Another possibility that makes sense in relation to the actual moves involved is the grouping (20, 10, 20, 10), (2H, 1H10, 2H), (10, 20, 10, 20). There is some suggestion in Fig. 4 supporting the idea that this grouping could have been developed by subjects after much practice.

The reason that error profiles and patterns of latency during learning apparently did not reflect conceptual structure may be a direct consequence of the tendency for this problem to be organized in a forward direction by subjects. The sequential tasks studied by other investigators involve learning series of digits (like 123543) and counting in base n . In both of these tasks, subjects find a simple pattern that generates a few digits, but that pattern is interrupted. (In the 123543 pattern, subjects will say "1234··," and the "4" is an error. In counting base four, the same kind of thing happens.) Errors at break-points in the structure being learned occur because subjects are applying a rule that has worked in past events. Thus, these tasks are organized retroactively, and the method of identifying structural units with difficult points during learning depends on that form of organization. The opposite kind of organization seems to characterize subjects' performance in this task.

REFERENCES

- AMAREL, S. On representations of problems of reasoning about actions. In D. Michie (Ed.), *Machine Intelligence* 3. Edinburg University Press, 1968.
- ATKINSON, R. C., BOWER, G. H., & CROTHERS, E. J. *An introduction to mathematical learning theory*. New York: Wiley, 1965.
- BJORCK, R. A. All-or-none subprocesses in the learning of complex sequences. *Journal of Mathematical Psychology*, 1968, 5, 182-195.
- ERNST, G. W., & NEWELL, A. *GPS: a case study in generality and problem solving*. New York: Academic Press, 1969.
- GREENO, J. G., & SCANDURA, J. M. All-or-none transfer based on verbally mediated concepts. *Journal of Mathematical Psychology*, 1966, 3, 388-411.
- HARLOW, H. F. Learning set and error factor theory. In S. Koch (Ed.), *Psychology: a study of a science*. Vol. 2. *General systematic formulations, learning, and special processes*. New York: McGraw-Hill, 1959. Pp. 492-537.
- NEWELL, A., & SIMON, H. A. *Human problem solving*. Englewood Cliffs, N. J.: Prentice-Hall, 1972.
- POLLIO, H. R., & REINHART, D. Rules and counting behavior. *Cognitive Psychology*, 1970, 1, 388-402.
- QUINLAN, J. R., & HUNT, E. B. A formal deductive problem-solving machine. *Journal of the Association for Computing Machinery*, 1968, 15, 625-646.
- REITMAN, W. R. *Cognition and thought*. New York: Wiley, 1965.
- RESTLE, F. The selection of strategies in cue learning. *Psychological Review*, 1962, 69, 329-343.
- RESTLE, F., & BROWN, E. Organization of serial pattern learning. In G. Bower (Ed.), *The psychology of learning and motivation: advances in research and theory*, Vol. 4. New York: Academic Press, 1970. Pp. 249-331.

- RESTLE, F., & GREENO, J. G. *Introduction to mathematical psychology*. Reading, Mass.: Addison-Wesley, 1970.
- SUPPES, P., & GINSBERG, R. A fundamental property of all-or-none models, binomial distribution of responses prior to conditioning, with application to concept formation in children. *Psychological Review*, 1963, **70**, 139-161.
- THOMAS, J. T., JR. An analysis of behavior in the hobbits-orcs problem. University of Michigan: Human Performance Center Technical Report No. 31, August, 1971.

(Accepted November 5, 1973)