NASA TECHNICAL
MEMORANDUM

NASA TM X-62,005

NASA TM X-62,005

PROBABILITY LEARNING: THE SHORTEST PATH HYPOTHESIS

by Edward M. Huff
Ames Research Center
Moffett Field, Calif. 94035
September 1970

PROBABILITY LEARNING: THE SHORTEST PATH HYPOTHESIS

Edward M. Huff

Ames Research Center

SUMMARY

In this study a specific hypothesis was tested concerning the development of a preference effect in human decision tasks that require predictions of future events. Six groups of subjects were exposed to different probabilistic sequences in which the recurrence paths to the preceding event varied in length. It was hypothesized that subjects would develop a preference for alternatives with the shortest recurrence paths. The results clearly support the hypothesis and show how the characteristics of the probabilistic environment influence human task performance.

INTRODUCTION

This study is part of a series that examines man's capability for predicting discrete events that occur with statistical rather than deterministic regularity. Of particular interest are the dependency properties that enhance or inhibit this ability.

In an earlier report (ref. 1) it was found that when a series of four stimulus events was generated by a homogeneous Markov process, the nature of the sub-sequences favored by the generator strongly influenced learning. In these generators any event could follow a prior event, but certain first-order transitions were favored by high probabilities. These high probabilities, in turn, predisposed the time series to contain certain dominant sub-sequences (e.g., runs of homogeneous events, event alternations, three-event cycles). When the subject's task was to predict which event would occur next, his ability was found to be inversely related to the length of the single-event cycle involved in the dominant sub-sequence.

One way to interpret the results of the previous experiment is to observe that subjects learned those event transitions most quickly that tended to *cluster* in the stimulus sequence. In fact, ease of learning was directly associated with the degree of negative skew in the recurrence time distribution to the prior stimulus. Although it is not clear why it should be so, this may have resulted in some learning advantage similar to a "massed practice" effect.

The present experiment carried this reasoning a step further, and directly tested whether the shorter of two equally probable recurrence paths would be preferred. The Markov generators depicted as flow diagrams in figure 1 were used. Here, transition matrices were constrained so that only two cells per row were nonzero, and each nonzero cell had a probability of 0.5. This type of matrix was used earlier (ref. 2) to examine "sequential guessing habits" but was selected for this study specifically because the subject has a choice on each trial between two equally probable alternatives. In the figure, the alternatives with the shorter recurrence paths are connected with bold arrows.

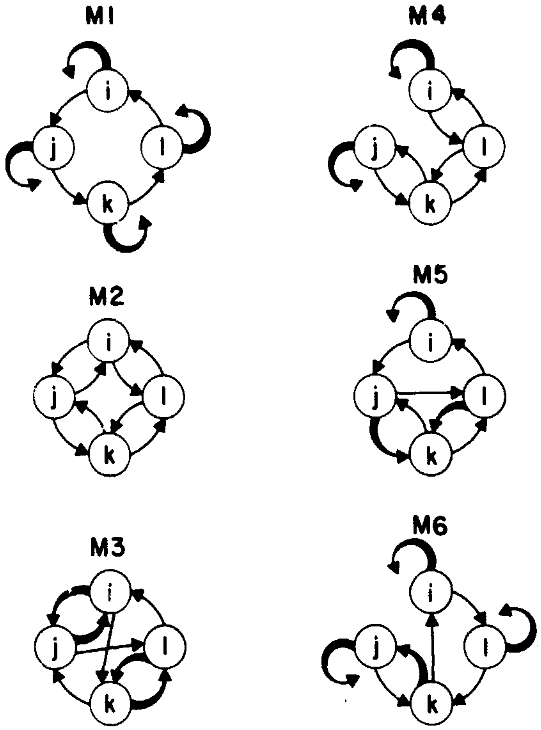

Figure 1.– Flow diagrams of Markov generators to control stimulus sequences. Heavy arrows show SP transitions where SP hypothesis applies.

The allowable transitions within each matrix determine how the events are "clustered" in the stimulus sequence. For example, in M1 each event is either followed by itself or some one other event. Note that once an event (e.g., E_i) does not recur, it is *impossible* for this event to appear again until all of the other events (i.e., E_j, E_k, and E_ℓ) have occurred at least once. In this case, each event tends to occur in homogeneous runs that are separated by at least three intervening events. In M6, on the other hand, three of the events can follow themselves, but the homogeneous runs tend to be distributed differently. Specifically, E_j runs may be separated by as few as one intervening event, namely E_k; but E_i and E_ℓ runs must be separated by at least two intervening events. Also note that in this generator, E_k is never followed by itself since one or more occurrences of E_j or E_k must intervene between its recurrences. The path through E_j however, will be associated with shorter recurrence times.

The experimental hypothesis to be tested was that subjects would *not* learn to choose each of the two equally probable alternatives on each trial with equal frequency. It was further hypothesized that a choice preference would develop that followed a simple rule: subjects should prefer to predict that event which has the shorter recurrence path to the preceding event. This will be called the SP (shortest path) hypothesis. Using M6 as an example, this means that having seen either E_i, E_j, E_ℓ the subject should tend to predict an event recurrence. Having just seen E_1, however, he should prefer to predict E_j rather than E_i, because that event provides the shorter path to E_k.

It may be noted that the SP hypothesis does not allow a differential prediction in all cases. In M2, for example, none of the available choices have shorter paths. In cases like this, it was expected that the frequency of choosing either alternative would be about equal.

METHOD

Experimental Groups

Seventy-two male university students between the ages of 17 and 27 were assigned at random to one of six experimental groups. For each group of 12 subjects, two sequences of 600 events were constructed that approximated the theoretical frequencies of one of the Markov processes in figure 1. Each test sequence was administered to half of the subjects in each group. For all sequences, marginal frequencies were within 4 percent and transition frequencies were within 5 percent of their theoretically expected values.

Design

As in the earlier experiment (ref. 1) subjects were required to predict which of four symbols (+, X, –, 0) would next appear on a viewing screen by pressing one of four buttons. Each of the buttons had one of the symbols inscribed above it. In order to avoid the possibility that symbol preferences would influence the results, the symbols were randomly identified with the events in the Markov matrix for each subject. A latin square procedure insured a balance of event-symbol assignments within groups.

Apparatus and Procedures

Two subject consoles were located in a large sound attenuated room and were separated by a wooden partition. Each console included a 5-inch Tektronix CRT, used to display stimulus events, and a set of five buttons located under the fingers of the subject's right hand. Each subject sat on a comfortable couch in a partially reclining position, with the CRT display located at a distance of approximately 18 inches along his horizontal line of sight.

Since the experiment was controlled by a LINC-8 computer, it was not necessary for the two subjects sitting at the consoles to be synchronized. Indeed, most often they belonged to different experimental groups. Each subject, therefore, received his instructions individually from the experimenter over a set of headphones.

After being seated, the subject was instructed to predict which one of the four symbols would next appear on his CRT display. A question mark was programmed to appear at the start of each trial to indicate when a prediction should be made. As soon as the subject pressed a button (except the thumb button, which was inoperative) the corresponding symbol was reflected below the question mark on the screen. A short time later the question mark was replaced by one of the four symbols, and the subject could compare his prediction with the correct symbol. For all experimental groups the response interval (question mark period) was 2.0 seconds, and each of the 600 stimuli in the sequence appeared on the screen for 2.5 seconds. A trial, therefore, took 4.5 seconds, and the total experimental session lasted 45 minutes. The instructions did not indicate that sequential dependencies could be found in the stimulus sequences. The apparatus did not allow the subject to change his prediction once a button was pressed.

RESULTS

The results of the experiment are summarized in the table. An *appropriate prediction* (AP) was considered to be either of the responses that had nonzero probability of being correct on a particular trial. This, of course, was determined by the matrix that controlled the stimulus sequence, and the symbol that had just appeared on the previous trial. Of the AP's, those that corresponded with the SP hypothesis are called SP (shortest path) predictions. Those that did not conform to the hypothesis are called LP (longest path) predictions.

In the table, the second column indicates, for the reader's convenience, the path lengths of the two appropriate alternatives following each event (see fig. 1). Cases where the SP hypothesis applies are shown with an asterisk. The next column shows the percentage of AP's that were made for each event in each group during the last 300 trials of the experiment. Since no significant differences were found in total AP's between subgroups that received different generator sequences, subgroup data were combined for purposes of this analysis. Overall, the high degree of AP learning is consistent with Bennett, Fitts, and Noble's (ref. 2) results, although a considerably greater range was obtained (76-97 percent). It is also conspicuous that the highest AP levels were obtained for those events where one of the appropriate alternatives was an event repetition (i.e., had a shortest path length of 1). In every case at least 90 percent was reached, a fact that was not true of any of the other events.

Since the SP hypothesis states that subjects should develop a preference for predicting the appropriate alternative that has the shortest path to the preceding event, it is merely necessary to partition the AP's into SP and LP predictions. The hypothesis asserts that SP predictions should exceed 50 percent of the AP's in all cases where the hypothesis applies.

Examination of the table shows that in *every* case more than 50 percent of the AP's were, indeed, SP predictions. Averaging over all cases, the percentage of SP predictions was 62 percent. In view of the fact that the hypothesis was not contradicted in a single instance, no further statistical tests are required.

One question that might be asked with regard to the preference effect is whether the amount of preference is related in some manner to the particular path lengths involved. Examination of the table shows, in general, that a greater effect is realized (i.e., 65 percent vs 57 percent) if the shortest path length is 1 rather than 2. It does not appear to be true, that the effect increases proportionately with the difference between the SP and LP lengths. This last conclusion is speculative, however, since the data do not provide a sufficiently large number of points for an overall comparison; to do so, the number of events that the subjects predict would have to be increased beyond four.

The procedure that was used to test the SP hypothesis could be criticized because the transition frequencies in the various sequences were only accurate to within 5 percent of their theoretically expected values. Unless some provision is made for this, then, the SP hypothesis could be confused with the interpretation that subjects learn small differences in the relative frequency of events and "maximize" their percentage of correct predictions. This will be called the *maximization hypothesis.* Indeed, examination of the SP transitions in the table (column 5) shows that there was a small overall bias in the stimulus sequences; that is, in six cases SP transitions exceeded 50 percent, and in only three cases were they less than 50 percent; overall, they occurred 50.5 percent of the time.

One way to circumvent this difficulty is to note that if subjects do learn small percentage differences in event transitions, and comply with the maximization hypothesis, then they should follow this policy consistently. In particular, they should employ it even in those cases where the SP hypothesis does not apply. These cases, then, can serve the useful purpose of allowing an independent evaluation of the maximization hypothesis.

In each case where the SP hypothesis did not apply, the appropriate alternative that occurred *most frequently* in the stimulus sequence was first identified. The last column in the table indicates the combined percentages of these transitions for the two test sequences from each generator. The percentage of AP's that were made of these most frequent (MF) events were then recorded, and are indicated in the next-to-last column. Again, data from subgroups that received different test sequences were combined. It is clear that under the maximization hypothesis each of these percentages should exceed 50 percent.

Since three of the seven cases in the table show response percentages below 50 percent, and the overall average for predicting the MF alternative was also slightly below this value (i.e., 49.97), no further statistical test is necessary. The data support the notion that appropriate alternatives with equivalent recurrence paths are predicted equally often, even though one alternative may occur slightly more frequently than the other. In short, the subjects do not discriminate small differences in event frequencies.

DISCUSSION

In the Bennett, Fitts, and Noble study (ref. 2), two 5-alternative Markovian stimulus generators were used, each of which stressed transitions that were either "concordant" or "discordant" with previously determined subject guessing habits. The authors reported that AP's were learned more quickly in the concordant sequence, presumably because the symbol transitions corresponded with the subject's normal guessing tendencies.

In the present experiment, the stimulus symbols were randomly identified with events in the Markov process for each subject, a fact that prevented *a priori* subject guessing habits from having systematic effects on group performance. In addition to the findings of Bennett *et al.*, therefore, it may also be concluded that guessing preferences result from the structural properties of the stimulus generator itself. Furthermore, it would appear, at least with the generators used in this experiment, that the preference effects can be adequately ascribed to the recurrence paths that the structure allows.

In recent years, a great deal of attention has focused on encodement procedures that subjects use in binary probability learning. The "run-length" hypothesis (ref. 3) posits that subjects encode sequences into numerical representations of successive run lengths. The "k-span" hypothesis (ref. 4) assumes that the subjects remember k units of the previous stimulus sequence which they encode as a single stimulus event. In both cases, the encoded stimuli are assumed to become associated with responses, and these, in turn, characterize the typical predictions made under various circumstances.

It is conspicuous that the run-length hypothesis, which accounts for a great deal of binary data, can account for very little of the present findings. Although it might apply to M1 sequences, where runs of homogeneous events were prevalent, it is not clear how it could explain the learning of the M2 or M3 sequences where runs greater than 1 could not occur. The hypothesis says nothing about how the subjects pick alternatives when an ongoing run discontinues, but merely capitalizes on the fact that binary sequences contain only one alternative.

It would seem at first that some elaboration of the k-span hypothesis, which is the more generic (although less specific) of the two, would be appropriate to explain multiple stimulus dependency learning. In the present experiment, nevertheless, only certain sub-sequences were created by any given generator. In each case, no more predictive information could be derived from the last k stimuli than from the most recent one, and it is easily verified (theoretically) that each prior sequence of k events occurred with equal relative frequency (i.e., there were no higher-order dependencies). There could be no advantage, then, in k being greater than 1, at least insofar as helping the subjects to learn first-order sequential dependencies.

The k-span hypothesis, then, merely allows for multiple stimulus dependency learning, but does not indicate the mechanism by which it takes place. With regard to the SP hypothesis, the k-span theory has only *post hoc* validity in that subjects did, indeed, learn certain sequences more thoroughly than others; but it does not predictively indicate which sequences should have been learned.

The major significance of the SP hypothesis, then, is that it is consistent with the earlier results (ref. 1), and that its verification helps to determine the necessary constraints for the construction of a model of temporal pattern learning. Two approaches seem particularly attractive at this time. First, it could be posited that the subject has a "short term store" mechanism with which to temporarily remember event transitions that have just occurred. Assuming that "long term" memory is modified as a function of the contents of the short term store, a number of model variations could favor the permanent retention of transitions that are multiply represented in the short term store. This, in turn, would result in a behavioral preference consistent with the SP hypothesis, because the shorter recurrence path would most often lead to multiple representation.

Second, if it is assumed that subjects attempt to remember past transitions from each event, but that their recollection deteriorates as a function of intervening trials, then a Bush and Mosteller learning mechanism (ref. 5) with an appropriately chosen decay operator could model the observed preference effect. Which of these model approaches will prove to be most valid will require further investigation.

Ames Research Center
National Aeronautics and Space Administration
Moffett Field, Calif., 94035, Sept. 17, 1970

REFERENCES

1. Huff, E. M.: Probability Learning: First-Order Markov Structures of Quarternary Events. NASA TN D-5684, 1970.

2. Bennett, W. F.; Fitts, P. M.; and Noble, M.: The Learning of Sequential Dependencies. J. Exptl. Psychol., vol. 48, 1954, pp. 303-312.

3. Rose, R. M.; and Vitz, P. C.: The Role of Runs of Events in Probability Learning. J. Exptl. Psychol., vol. 72, 1966, pp. 751-760.

4. Vitz, P. C.: Information, Run Structure and Binary Pattern Complexity. Perception and Psychophysics, vol. 3(4A), 1968, pp. 275-280.

5. Bush, R. R.; and Mosteller, F.: Stochastic Models for Learning. Wiley, New York, 1956.

TABLE 1.– PERCENTAGES OF APPROPRIATE (AP), SHORTEST PATH (SP), AND MOST-FREQUENT (MF) PREDICTIONS FOR INDIVIDUAL EVENTS IN THE SIX MARKOV STRUCTURES ON THE LAST 300 TRIALS

Structure		Path length	Percent AP's	SP hypothesis		Maximization hypothesis	
				Percent SP predictions	Percent SP transitions	Percent MF predictions	Percent MF transitions
M1	E_i	1*-4	95	66	50		
	E_j	1*-4	95	63	50		
	E_k	1*-4	97	62	51		
	E_ℓ	1*-4	93	62	50		
M2	E_i	2-2	88			52	52
	E_j	2-2	87			52	50
	E_k	2-2	87			46	50
	E_ℓ	2-2	89			48	51
M3	E_i	2*-3	76	53	51		
	E_j	2*-3	78	54	50		
	E_k	2*-3	78	55	52		
	E_ℓ	2*-3	83	58	50		
M4	E_i	1*-2	92	63	52		
	E_j	1*-2	91	65	49		
	E_k	2-2	89			49	50
	E_ℓ	2-2	88			53	51
M5	E_i	1*-3	94	69	50		
	E_j	2*-3	84	55	48		
	E_k	2-2	84			50	53
	E_ℓ	2*-3	83	67	52		
M6	E_i	1*-3	90	66	50		
	E_j	1*-2	90	70	49		
	E_k	2*-3	84	56	50		
	E_ℓ	1*-3	92	65	52		

NOTE: Percent SP transitions and percent MF transitions refer to the actual percentage of shortest path and most-frequent transitions that occurred in the stimulus sequences, respectively.