Multivariate analyses of performance in a DRL paradigm

FRED H. GAGE and SELBY H. EVANS
Texas Christian University, Fort Worth, Texas 76129

and

DAVID S. OLTON
The Johns Hopkins University, Baltimore, Maryland 21218

Multivariate analyses were used to explore the interrelations between response measures used in a DRL-20 schedule. A principal components analysis revealed two major components which accounted for a majority of the variance in all measures. One component included responses, reinforcements, efficiency, and a decrease in short-latency interresponse times; it could be viewed as supporting the previous use of the construct inhibition in the interpretation of DRL results. The other component was primarily associated with increased responding in the first period in which reinforcement was available; it could be viewed as a measure of temporal discrimination. A discriminant function, with blocks of trials as the criterion, was used to obtain from IRT/OP data a single measure summarizing change in behavior as a function of experience. The discriminant function incorporated both of the above sources of variance. The findings indicate that currently used summary measures do not fully reflect changes that occur under the DRL schedule.

Although the differential-reinforcement-of-low-response-rate schedule (DRL) is often used to study animal behavior, results obtained with this schedule are reported in a variety of ways. Four different dependent measures are available for describing DRL performance: (1) response rate—the number of responses per unit time, (2) reinforcement rate—the number of reinforcements per unit time, (3) efficiency—usually the response rate divided by the reinforcement rate divided by the response rate, (4) IRT/OP—the conditional probability suggested by Anger (1956), which estimates the probability of responding in any particular time interval given an opportunity to respond in that interval.

The DRL procedure has received considerable attention, not only as a source of interesting data in its own right (Kramer & Rilling, 1970), but also as a means of assessing various neuroanatomical treatments (Johnson, Olton, Gage, & Jenko, 1977) and chemical treatments (Ellen, Aitken, Sims, & Stahl, 1975). Of particular importance to the present paper are two alternative explanations of DRL performance. The first treats successful DRL performance as a measure of temporal discrimination (Ellen & Aitkin, 1971; Ellen, Wilson, & Powell, 1964). According to this view, DRL reflects timing behavior and indicates the ability of the animal to determine the passage of time since the last barpress.

This research was supported in part by NIMH Grant 1-RO1-MH24213-01 to D.S.O. and TCU/Research Foundation Grant to F.H.G.

The second interpretation of DRL performance sees it as a measure of response inhibition. In this case, "inhibition" is used in a very descriptive fashion and refers to the ability of the animal to inhibit or prevent a response before the timing interval has ended. Several questions are posed as a result of these uses. Specifically, can data obtained from the DRL paradigm represent both inhibition and temporal discrimination at the same time? Can such constructs be represented by several different measures equally well and with the same implications? The purpose of this study is to illustrate methods of inquiry and present results relevant to these questions and to other questions which follow from them.

METHOD

Subjects

The subjects were 34 male Sprague-Dawley (Camm) albino rats, 90-120 days old at the beginning of the experiment. All animals were housed individually with ad-lib water.

Apparatus

All behavioral testing was carried out in two operant chambers contained in sound-attenuating compartments. In the middle of one wall of each operant chamber was located a lever, with a food cup 5 cm to the left of the lever and a water spout 7 cm to the right of the lever. Data were collected on digital counters which recorded the number of leverpresses occurring within the first 12 successive 4-sec intervals after initiation of the timing interval. A single 45-mg Noyes pellet was used as a reinforcer for a correct response.

Behavioral Testing

All rats were food deprived to 85% of their ad-lib body weights

by restricting food intake over a period of 5 days. They were maintained at this level throughout behavioral testing by being provided with the appropriate amount of food in their home cages approximately 1 h after each day's test. All behavioral test sessions were 30 min long. For the first 2 days of testing, leverpresses were shaped by food powder on the bar; reinforcement was provided on a CRF schedule. On the 3rd day, the DRL 20-sec schedule with unlimited hold was begun. Testing continued daily for a total of 30 days under the DRL schedule. Further details are presented in Johnson, Olton, Gage, and Jenko (1977).

Data Analysis

Performance was analyzed in six blocks of 5 days each. Four different behavioral measures, as suggested by Kramer and Rilling (1970) were obtained: (1) average number of daily reinforcements, (2) average number of daily responses, (3) average daily efficiency (number of reinforced responses/total number of responses), and (4) interresponse times/opportunity (IRT/OP). The first three behavioral measures were determined by obtaining the median value for the indicated period for each rat and then taking a mean of these medians. The IRT/OP measure was obtained from the formula discussed by Anger (1956). Twelve counting bins of 4 sec were used, but the results for only the first seven were retained for analysis.

The IRT/OP formula uses a denominator which is determined by counts in other bins. Thus the IRT/OP measure would tend to introduce correlations between intervals simply because each value is based on data from several intervals. Such an effect would be small in the early intervals. It would become increasingly important in the later intervals because the denominator is formed from fewer and fewer intervals. The reliability of the measure would also suffer in the later intervals because the denominator is based on relatively few counts.

To some extent, these problems are reduced by dropping the last few intervals, but the IRT/OP measure is not necessarily the best choice for multivariate analysis. We use it here to maintain continuity with previous research.

RESULTS

The data were analyzed in several ways, as dictated by particular research questions. The analyses should be considered exploratory rather than confirmatory in the sense characterized by Tukey (1977).

(1) Do the summary measures, reinforcements and

efficiency, represent the same behavioral changes, and do they adequately represent or summarize the IRT/OP data? In other words, can the measures be used interchangeably to represent DRL behavior change?

Table 1 presents the intercorrelations among reinforcements, responses, efficiency, and the IRT/OP counts. The correlations were formed by cumulating over both subjects and blocks of trials. Thus they represent both covariation attributable to subjects and covariation attributable to blocks. A dummy variable, trial block, taking on values 1, 2 . . . 6, is included in the table to indicate the covariation directly associated with blocks.

Table 1 shows a correlation of .94 between reinforcements and efficiency. In these data, at least, the two measures clearly represent the same phenomenon and could be used interchangeably. The intercorrelation table also shows that these variables have substantial correlations with some of the IRT/OP periods but the table does not make clear whether the IRT/OP data are adequately represented by these variables. Principal components analysis was used to inquire further into that question.

Principal components analysis is a method for analyzing sets of intercorrelated data so as to extract redundant aspects. It is related to factor analysis and sometimes is used in that context. The present use, however, is for the purpose of description and summary; it should not be confused with efforts to find factors of theoretical import. The results of principal components analyses are usually not regarded as representing theoretically interpretable factors and are normally subjected to rotations such as Varimax (Kaiser, 1958) to produce more interpretable factors.

The mathematical theory of principal components analysis is presented quite well by Tatsuoka (1971). Programs are available in most multivariate statistical packages.

Table 1
Intercorrelations Among Variables of Interest in Summarizing DRL Data

	BLK	RNF	RSP	EFF	IR1	IR2	IR3	IR4	IR5	IR6	IR7
BLK											
RNF	.47										
RSP	-.52	-.63									
EFF	.47	.94	-.73								
IR1	-.40	-.22	.78	-.41							
IR2	-.70	-.64	.83	-.67	.64						
IR3	-.55	-.82	.63	-.76	.21	.77					
IR4	-.32	-.90	.46	-.82		.50	.82				
IR5		-.73	.22	-.68			.42	.73			
IR6	.39	-.16	-.16	-.18	-.26	-.26		.26	.57		
IR7	.37	.27	.27	.26	-.24	-.30	-.16			.50	
D1	-.73	-.43	-.43	-.45	.60	.93	.60	.26		-.56	-.49

Note—Correlations are taken over subjects and blocks, $n = 150$. Correlations less than .15 in magnitude are omitted. Abbreviations are as follows: BLK, block number; RNF, average number of daily reinforcements; RSP, average number of daily responses; EFF, number of reinforced responses divided by total number of responses; IR1-7, the IRT/OP measure in the first seven 4-sec periods of the DRL schedule; D1, score on the first discriminant function based on differences between blocks of trials.

An intuitive understanding of principal components analysis can be developed as follows: Consider a scatter plot of scores on two variables expressed as standard scores and showing substantial intercorrelation. The plot will form an ellipse centered at the origin and tilted at roughly a 45° angle. The same data could be represented by a new set of dimensions lying along the major and minor axes of the ellipse.

In terms of this new set of dimensions, the data would show no correlation. The covariation between the two original variables would have been converted into variance along the major axis of the ellipse. Placement on this axis would account for a substantial proportion of the variance in each of the original variables. This change of axes is approximately what principal component analysis does.

The change of axes is accomplished by forming new variables (axes) which are weighted sums of the original axes (the original variables). Thus, another way of describing principal components analysis is that it forms new variables which are weighted sums of the original variables. The new variables are statistically independent of each other. Furthermore, the first variable (factor) accounts for the largest amount of variance possible, and each successive variate accounts for the largest amount of variance possible given the preceding variables.

In some cases, it is of interest actually to form the new variables, but for present purposes, the primary interest is in how well the factors represent the original variables. That information is best expressed in the correlations between the original variables and the factors.

Table 2 presents a principal components analysis of the variables in question. Table 2 shows that the first factor accounts for more than 50% of the variance in the data and that this factor does, indeed, represent reinforcements and efficiency. Response rate also correlates well with this factor, but not as well as the other two summary measures. IRT 2, IRT 3, and IRT 4 also are strongly represented by this factor. Thus we can conclude that reinforcement rate and efficiency represent a major component of the IRT/OP data. In addition, a second factor accounts for nearly 25% of the variance and is the main representative of IRT 6, the optimal period for responding under the DRL schedule. Thus, it seems that not all the variation in the IRT/OP data is adequately represented by the two summary variables, reinforcement rate and efficiency.

(2) Is there a best single measure to summarize the IRT/OP data?

The first part of the answer is foreshadowed by the previous analysis. Principal components analysis of the IRT data by itself produced a pattern similar to that in Table 2. But the first factor accounts for 31% of the variance. Thus, it seems that at least two mea-

Table 2
Correlations of Original Variables with the First Two Factors from Principal Components Analysis of Variables of Interest in Summarizing DRL Data

	F1	F2
RNF	.94	.21
RSP	-.82	.42
EFF	.95	
IR1	-.46	.64
IR2	-.81	.43
IR3	-.87	
IR4	-.84	-.40
IR5	-.61	-.67
IR6		-.82
IR7	.30	-.53
PROVAR	.53	.24
CUMVAR	.53	.77

Note—Correlation matrix is given in Table 1, except that D1 was not used in this analysis. Correlations less than .15 in magnitude are omitted. Row abbreviations are as follows: RNF, average number of daily reinforcements; RSP, average number of daily responses; EFF, number of reinforced responses divided by the total number of responses; IR1-7, the IRT/OP measure in the first 4-sec periods of the DRL schedule. The row marked PROVAR gives the proportion of variance accounted for by each factor. The CUMVAR row gives cumulative proportion of variance accounted for.

asures, one representing each of the two factors, are needed to summarize the IRT/OP data.

If one must choose a single measure, presumably Factor 1, which accounts for the greatest percentage of variance, is a good candidate. But the question is really ill defined unless we specify "best" in terms of some criterion. If the DRL paradigm is viewed as a learning task, perhaps the criterion should be that the measure best represents or measures learning.

The theoretical and conceptual problems in the definition of learning can be avoided if we assume that learning, in this context, is the change in measured behavior, as a function of trials (really, blocks of trials) in the DRL schedule.

Under that view, the most direct approach would be to ask what changes in IRT/OP behavior can be written as a function of (predicted from) blocks of trials. That approach is not as direct as it sounds, of course, because we would still have to select the form of the equation—linear, quadratic, or perhaps some particular learning model. In the present case, however, we have seven measures of behavior which vary as a function of trials. What we seek is a summary measure derived from them which best expresses their variation as a function of trials. To obtain this, we can reverse the prediction—that is, seek a function that "predicts" blocks of trials from the IRT/OP data.

What function? If we accepted certain theoretical positions, we might be led to choose a linear or at least a monotonic relationship. If we eschew theoretical positions, however, we can seek simply that func-

tion of the IRT/OP which maximally discriminates between blocks of trials—that function which maximizes the between-block variance relative to the within-block (between-subjects) variance. In other words, the discriminant function.

Discriminant function analysis is a procedure designed to reduce multiple, correlated measures to one or a few independent variables which are best suited to discriminate between predefined groups. The mathematical basis is presented by Tatsuoka (1971) and other texts on multivariate analysis. Programs are available in most statistical packages. The procedure finds a set of weights for the original measures such that the weighted sum maximizes the between-group variance relative to the within-group variance. The result is termed the first discriminant function or variate. Additional variates are similarly defined, with the restriction that they are orthogonal to those previously established. The weights for the discriminant function are of interest in that their magnitude indicates the importance of the variables in forming the function.

A discriminant function analysis was done using the seven IRT/OP variables to discriminate among the six blocks of trials. The weights for forming the first discriminant variate were 4.6 for IRT 2, $-.82$ for IRT 6, and relatively small negative values for the remaining periods. Thus, both of the factors identified earlier are needed to define learning in the sense used here. This variate could be regarded as strongly representing between-group differences, since 76% of its variance is associated with differences between groups.

IRT Period 2 dominates this variate. In considering this result, one should keep in mind that the discriminant function maximizes the ratio of between-group variance to within-group variance. Thus, its heavy weighting of IRT 2 does not necessarily imply that the other IRT means are less strongly related to trials. Rather, the low weights given to the other IRT periods may indicate that they show greater within-group variation. Other examinations of the data suggest that this is the case.

In any event, the discriminant function analysis indicates that a single variate based on the IRT/OP data can represent most of the change in that data as a function of trials. Moreover, that change is best characterized as a decrease in the frequency of the relatively short-latency responses which fall in IRT Periods 2, 3, and 4. This observation could be viewed as supporting the previous use of the construct inhibition in the interpretation of DRL results.

The correlations between the resulting discriminant variate and the other variables of interest are given as the last row (DI) in Table 1. The variate does not correlate well with either reinforcement rate or efficiency. It does have a substantially greater linear cor-

relation with blocks than does either of these variates. Thus it appears that the discriminant variate may be a better summary measure than reinforcement rate and efficiency, if the objective is to summarize in one measure the change in behavior as a function of trials.

(3) Can all subjects be summarized by a single measure?

Since the above analyses have been performed on data which include the variance associated with individual differences, it is of interest to ask whether the two factors suggested above might result from differences among individuals in the way they responded to the task. Accordingly, a Q factor analysis was conducted to investigate the possibility of clusters of subjects showing different patterns.

The Q factor analysis was accomplished by reversing the roles of subjects and measures as those were treated in the previous factor analysis. In other words, the analysis was done on an intercorrelation matrix containing correlations between subjects, rather than between measures. The correlations were based on the responses of the subjects in each IRT period in each block of trials. This arrangement provided only 42 observations for each of the 24 subjects. While a much larger number of accumulated observations would have been desirable, it was not available. The results of the analysis indicated that the first factor accounted for 69% of the variance, with the remaining factors accounting for 9%, 6%, and 3%, respectively. This result indicates that the subjects were generally similar in the behavioral pattern over the task.

(4) Is it possible to derive measures of both response inhibition and temporal discrimination from the IRT/OP data?

Answering this question requires some assumptions about what those expressions mean in terms of empirical referents. The most direct interpretation seems to us to be that inhibition is represented by an absolute decrease in the frequency of short-latency responses and discrimination is represented by a relative increase in the frequency of responses in the first reinforced interval. Obviously, in this sense, empirical discrimination can only be displayed if there is empirical inhibition. But the reverse is not true, so that there is at least a logical possibility of observing separate variance associated with empirical discrimination.

The principal components analyses of the IRT data suggest the empirical possibility of such observations. IRT Periods 2, 3, and 4 correlate strongly with Factor 1, while Period 6 correlates strongly with Factor 2. It seems reasonable to view Factor 1 as representing the suppression of short-latency responses, and unavoidable to view Factor 2 as the emission of responses at the optimal time. If so, then the standard measures, reinforcement rate and efficiency, are primarily

Figure 1. Means and 95% confidence intervals for the IRT/OP proportions for IRT 3 and IRT 6 as a function of blocks. Also, means and 95% confidence intervals for reinforcements and D1 as a function of blocks. Reinforcements has been expressed as a proportion of the estimated maximum to put it on the same scale as the IRT/OP measure. The scale for D1 is arbitrary, since it is a derived measure; D1 has been rescaled so that its units are comparable to the other plots. Note that RNF and IRT 3 are similar in shape and essentially linear. IRT 6 and D1 are curvilinear and fairly similar in shape. The graphs illustrate that rather different conclusions about the progress of learning might be reached, depending on what portion of the data was considered.

measuring empirical inhibition—at least according to the testimony of Table 2, in which reinforcement rate and efficiency correlate highly with Factor 1.

Figure 1 presents plots of key variables relevant to this issue. Reinforcements can be seen as clearly a reflection of IRT 3. Both show nearly linear changes. IRT 6 shows a curvilinear change evidenced in particular by the clear separation of the confidence interval for Block 1 from those for the remaining blocks. This pattern is not well represented by reinforcements, but the discriminant function variate does capture and express this pattern.

DISCUSSION

It appears that different ways of treating the DRL data give somewhat different pictures of the behavioral changes taking place. It is possible that researchers studying inhibition or discrimination might wish to select the form of data treatment that emphasizes the phenomenon in which they are interested. It should be noted that the relatively modest differences between IRT 3 and IRT 6 in the present context

might be greatly amplified by treatments designed to modify one or the other of these inferred processes. In such cases, the selection of the appropriate variable could be crucial.

Inspection of the data indicated that most of the change in performance occurred during the first 2 weeks (the first two blocks of trials). Thus the data may not have provided sufficient temporal resolution to permit a full display of possible differences in the course of development of inhibition and temporal discrimination.

In conclusion, these analyses indicate that the summary measures, reinforcement rate and efficiency, represent essentially the same variance as found in the data. These two measures, however, do not capture all of the meaningful variation in the IRT/OP data. Production of high rates during the earliest period in which reinforcement is available may be a source of variance somewhat distinct from the suppression of short-latency responses, indicating that two processes may be taking place during the acquisition of DRL. These results also suggest that multivariate techniques may provide useful insights in the examination of such data.

REFERENCES

- ANGER, D. Dependence of interresponse times upon the relative reinforcement of different interresponse times. *Journal of Experimental Psychology*, 1956, **52**, 145-161.
- ELLEN, P., & AITKIN, W. C. Absence of temporal discrimination following septal lesion. *Psychonomic Science*, 1971, **22**, 129-131.
- ELLEN, P., AITKIN, W. C., JR., SIMS, T., & STAHL, J. M. Cholinergic blockade, septal lesions, and DRL performance in the rat. *Journal of Comparative and Physiological Psychology*, 1975, **5**, 409-420.
- ELLEN, P., WILSON, S. S., & POWELL, E. W. Septal inhibition and timing behavior in the rat. *Experimental Neurology*, 1964, **10**, 120-132.
- JOHNSON, C. T., OLTON, D. S., GAGE, F., III, & JENKO, P. S. Hippocampus, hippocampal connections, and DRL. *Journal of Comparative and Physiological Psychology*, 1977, **91**, 508-522.
- KAISER, H. F. The varimax criterion for analytic rotation in factor analysis. *Psychometrika*, 1958, **23**, 187-200.
- KRAMER, T. J., & RILLING, N. Differential reinforcement of low rates: A selective critique. *Psychological Bulletin*, 1970, **74**, 225-254.
- TATSUOKA, M. M. *Multivariate analysis: Techniques for educational and psychological research*. New York: Wiley, 1971.
- TUKEY, J. W. *Exploratory data analysis*. Reading, Mass: Addison-Wesley, 1977.

(Received for publication May 2, 1978;
revision accepted January 4, 1979.)