ORIGINAL ARTICLE

S. M. Hartz · Y. Ben-Shahar · M. Tyler

Logistic growth curve analysis in associative learning data

Abstract We propose an alternative statistical method, logistic growth curve analysis, for the analysis of associative learning data with two or more comparison groups. Logistic growth curve analysis is more sensitive and easier to interpret than previously published methods such as χ^2 or ANOVA, which require the data to be collapsed into individual total scores or proportion of responses over time. Additionally, this type of analysis better fits the typical graphical representation of associative learning data. An analysis is presented where associative learning data from honeybees are analyzed using the three techniques, and the accessibility and power of the logistic growth curve analysis is highlighted.

Keywords Logistic growth curve analysis · Comparison of statistical methods · Associative learning · Honeybees

Introduction

The goal of many learning and memory experiments is to identify and characterize trends in cognitive skills for groups of individuals. One of the commonest paradigms used in such studies is classical conditioning (reviewed by Menzel 1999). Classical conditioning studies use a repeated exposure to a conditioned stimulus (CS) that originally had no obvious biological significance to the tested subjects, coupled with a biologically relevant unconditioned stimulus (US). Repeated exposures to the coupled CS-US result in increased performance on the tested learning criterion (usually an acquired behavioral response to the CS).

The data from these studies are a sequence of responses and non-responses for each animal. In previous reports, some authors scored the number of responses and did a *t*-test to evaluate the difference between groups; other authors calculated proportions of responses at each time point for each group and did a χ^2-test on the difference between groups at each time point. Both methods address some aspects of the phenomenon, but ignore others. We advocate a third approach, logistic growth curve analysis, and argue its advantages. To illustrate the theory, we analyzed a single data set using the three different methods for comparative purposes.

S. M. Hartz (✉) · M. Tyler
Department of Statistics,
University of Illinois at Urbana-Champaign, 101 Illini Hall,
725 S. Wright St., Champaign, IL 61820, USA
e-mail: hartz@stat.uiuc.edu,
Tel.: +1-217-2447198, Fax: +1-217-2447190

Y. Ben-Shahar
Department of Entomology,
University of Illinois at Urbana-Champaign,
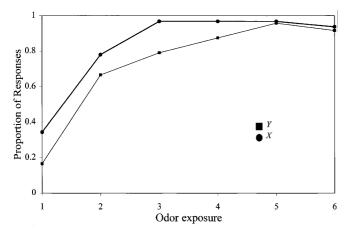Champaign, IL 61820, USA

Methods

Data set

Ben-Shahar et al. (in press) trained 56 honey bees to associate odor A (geraniol) with a sugar reward in a sequence of six odor exposures. If the bees learned to associate odor A with a reward, the proportion of bees responding would increase with time. The resulting data for each bee is a sequence of responses (R) and non-responses (N). For example, the response sequence of a bee to odor A of N-R-N-R-R-R would suggest that the bee had learned the association.

Bees were collected from two different behavioral groups, group X (*n*=32), and group Y (*n*=24), to determine whether they perform differently in this learning assay. The raw proportion of responses to odor A is plotted in Fig. 1. Visually, there appears to be a trend in the plots: bees from group X learn more quickly than bees from group Y. Initially, they respond more frequently, and reach the maximum level of learning by the third learning trial. The bees from group Y do not catch up with the bees from group X until the fifth learning trial.

The three analytic methods

Method 1: t-test on scored data

The data can be analyzed by collapsing the data to form either one observation per subject or one observation per time point. In order

Fig. 1 Learning curves. Proportion of trials with responses to an odor by 32 bees in group X and 32 bees in group *Y*. Logistic growth curve analysis uses a logistic curve (sigmoid curve between zero and one) to model these two curves and statistically compare them, after adjusting for positive correlation between observations from the same bee

to create one observation per subject, the data from each subject were converted to the total number of responses (e.g., Bitterman et al. 1983; Bhagavan et al. 1994; Pelz et al.1996; Gerber et al. 1998).

Information is lost, however, when the data are collapsed in this way. For example, in an experiment with six trials, both the response sequences R-N-R-N-R-N and N-N-N-R-R-R give a score of three responses. The first case appears to be a random fluctuation, but the second case is suggestive of learning. In fact, in an experiment with six trials, there are 20 different ways that a subject can achieve a total of three responses.

A *t*-test of the difference in total number of responses per bee in groups X and Y was performed. Both X and Y groups had bees whose response totals ranged from zero to six.

Method 2: χ^2-test on proportional data

These data can also be analyzed by converting the data to proportions of responses per group in each odor exposure (e.g., Ray and Ferneyhough 1997, 1999). Proportions were analyzed by using χ^2-tests (with 1 *df*) for each odor exposure separately. A Fisher's exact test was used when an entry in the table had $n \leq 5$.

Here, again, information is lost. Suppose there is an experiment with ten animals. Two out of ten animals respond initially, six different animals respond on the last trial, and there were random fluctuations in the middle. Overall, it appears the group of animals learned to respond because the proportion of responses increased. But there is no evidence of individual learning because different animals responded at different times. Additionally, there is difficulty in interpreting learning dynamics as a time-dependent process since each time point is analyzed as a discrete unit.

Method 3: logistic growth curve analysis

A common and intuitive way of presenting learning and memory data is to form „learning curves" by plotting the proportion of responses at each trial for a group of subjects as in Fig. 1 (e.g., Bitterman et al. 1983; Bhagavan et al. 1994; Gerber et al. 1996; Pelz et al. 1996). The learning curves from several experimental groups are frequently plotted on a single graph to allow visual comparison. Although the papers cited here presented learning curves, none of them used all the data from the curves to test the magnitude of the difference between the tested groups. Logistic growth curve analysis precisely tests the difference between several learning curves. It retains the structure of the responses, both
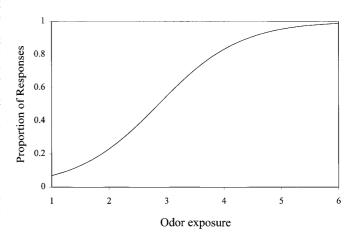
on the individual and group level, and results in a statistical comparison of predicted curves resembling the empirical learning curves.

The data addressed in this paper (Fig. 1) consist of repeated measures of the same subject, where a measurement is a response or a non-response (coded 1 for response and 0 for non-response). In an attempt to estimate the curves seen in Fig. 1, the inclination may be to fit a line or curve to the data using linear regression. But a linear regression of these data violates several assumptions of the linear regression model. First, the proportions of responses to be modeled fall only between zero and one (equivalently, between 0 and 100%), so the assumption that a line or a polynomial can fit the data is incorrect because anything but a flat line will eventually cross either one or zero. Secondly, knowing the response of a subject at one trial gives a better idea of what its response will be at a later trial. That is, each consecutive response of a single subject is dependent on previous responses. For instance, if a subject were above average, we would expect it to stay above average. Statistically speaking, the responses of a single subject are correlated.

Logistic growth curve analysis, the combination of logistic regression and growth curve analysis, incorporates non-linearity and correlation into the analysis. Logistic regression models the probability of response over time, and its sigmoid curve lies between zero and one (Fig. 2). Similarly to linear regression, it contains a slope parameter and an intercept parameter. The ratio of the intercept to the slope marks the point in time where half the subjects have responded. The slope parameter indicates the steepness of the curve. A slope between zero and one flattens the curve, and a slope of zero indicates no predicted change over time.

Growth curve analysis assigns a correlation structure to the data, allowing a positive correlation for observations from the same individual, while maintaining the observations from different individuals as independent. This correlation structure means that observations from a single subject are more related to each other than they are to observations from a different subject. The correlation estimate is then used to calculate the logistic regression parameters and their standard errors (Dunlop 1994; Littell et al. 1996).

We performed logistic growth curve analysis using PROC GENMOD in SAS, version 6.11 (Littell et al. 1996; see Appendix). A logistic curve fit to the data and then the correlation structure for each individual bee was estimated with generalized estimating equations, or GEE (for more information on correlation structure estimates see Littell et al. 1996, SAS 6.12 manual). Finally, the cor-

Fig. 2 Typical logistic curve plot. This curve serves as a model for learning: an animal initially does not respond, and gradually becomes highly responsive. The formula for the curve is $f(x)=1/(e^{mx+b}+1)$. The *point at which the curve turns upward* corresponds to the intercept of the curve (*b*), and the *rate with which the curve increases* after it turns corresponds to the slope of the curve (*m*)

relation estimate was used to adjust the logistic regression estimates.

Performance in this assay is indicated by a combination of the intercept and the slope of the logistic curve. Because learning is almost immediate and equal in all groups tested (there is almost always at least one bee responding on the second odor exposure in each group), the intercept parameter should be approximately the same for all the groups. We therefore set the intercept to be the same in both groups. This allows the difference in slope between the two groups to represent the difference in learning between them. When the difference in slope is statistically significant, the difference between learning rates is statistically significant.
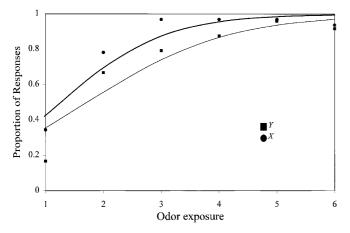
Results

The *t*-test calculates a mean score of 4.97 for group X (SD 1.20) and a mean score of 4.38 for group Y (SD 1.56). The *p*-value for a significant difference between the two groups is not significant at 0.11. The six χ^2-test results presented in Table 1 are harder to interpret. At each odor exposure, group Y has a lower proportion of responses than group X, but none of these are significant, so no statistically valid conclusions can be drawn.

The results of the logistic regression are seen in Fig. 3 and Table 2. The parameters of the estimates construct a line, modeling the logit of the proportion of bees that have

Table 1 Results from method 2: χ^2 analysis on proportional data (each odor exposure was analyzed separately)

Odor exposure	% Response		
	Group X	Group Y	P
1	34.4	16.7	0.24
2	78.1	66.7	0.51
3	96.9	79.2	0.07
4	96.9	87.5	0.30
5	96.9	95.8	1.00
6	93.8	91.7	1.00

Fig. 3 A fit of regression curves to learning data. Proportion of trials with responses to an odor and the estimated logistic regression growth curves. The *curves* are significantly different from one another (*P*=0.03)

Table 2 Results from method 3: logistic growth curve analysis

Parameter	Estimate	SE	P
Intercept	–1.50	0.63	0.02
Trial	1.51	0.41	0.00
Trial×behavior (learning parameter)	–0.62	0.29	0.03

learned (the log of the proportion that have learned minus the log of the proportion that have not learned). The intercept parameter corresponds to the point at which the logistic curve turns upward, the trial parameter corresponds to the rate at which the bees from group X learn, and the trial×forager parameter corresponds to the difference between the rate of learning in both groups. The P values are the probabilities that the corresponding parameters are actually zero. The significant negative parameter trial×behavior (*P*=0.03) indicates that bees from group Y learn significantly more slowly than bees from group X in this data set (α=0.05).

The adequacy of the logistic model was determined by comparing the estimated to the observed proportions, using the standard error of the estimated proportions $\sqrt{p(1-p)}$, where *p* is the observed proportions). Except for odor exposure 6, the observed proportions predominately fall within 1 SE of the predicted values and all fall in the standard 95% confidence interval around the predicted values. Thus the model is accepted as adequate.

Discussion

The above example demonstrates the power of logistic growth curve analysis to detect differences in learning curves as compared to a *t*-test or χ^2-test. Situations exist where the *t*-test or χ^2-test is significant and the logistic growth curve is not. But difference in total number of responses might not reflect differences in overall trends. If the scores from the two groups fluctuate between time points, an analysis may show significant differences between the two groups at one or more of the time points (by using χ^2 analysis), without a difference in overall trend. Since associative learning is a process over time, looking at a trend is a better reflection of this process than a total score or a proportion of correct responses.

Another advantage of the logistic growth curve analysis is that it is more flexible. Learning can be defined in several ways by fixing either the intercept or the slope parameter. More parameters can be added if more than one effect is researched, or confounding variables are suspected.

Logistic growth curve analysis is not the only statistically sound approach to analyzing associative learning data. Any statistical technique that both adjusts for the inherent correlation structure and accounts for both the dichotomous data and the pattern of learning over time is appropriate. Essentially, this includes „growth curve" analyses with asymptotic, increasing or decreasing func-

188

tions. The reason we chose to present logistic growth curve analysis is because it is available in many statistical packages without having to use any sophisticated programming.

Studies using this analysis can draw strong and precise conclusions concerning differences in learning performance between tested groups by comparing statistical parameters that correspond to learning rates. The technique is a statistical extension of a standard visual presentation of learning performace data.

Appendix

SAS code for logistic growth curve analysis

This program is to run a logistic growth curve analysis on the difference between two groups: group X and group Y. The data set contains the following variables:

1. num – identification number for each individual. These do not need to be sequential. They do, however, need to be numeric.
2. response – takes on values 0 or 1 representing whether the animal responded at that trial.
3. trial – trial number.
4. groupX – takes on values 0 or 1 indicating membership in group X.

Option one for data entry

```
data new;
input num response trial group;
cards;
*Insert data with the variables ordered in the above order.
Variables should be separated by spaces. In this case, the
responses for a single animal must be entered in separate
lines, according to the trial;
run;
```

Option two for data entry

```
data first;
input num colony respons1 respons2 respons3 respons4
respons5 respons6 groupX;
cards;
*Insert data with the variables ordered in the above order,
variables should be separated by spaces. Here, re-
spons1–response6 represent the responses for animal
"num" in trials 1–6 respectively;
run;
data one;
set first;
```

```
response=respons1;
trial=1;
keep num response trial groupX colony;
run;
data two;
set first;
response=respons2;
trial=2;
keep num response trial groupX colony;
run;
data three;
set first;
response=respons3;
trial=3;
keep num response trial groupX colony;
run;
data four;
set first;
response=respons4;
trial=4;
keep num response trial groupX colony;
run;
data five;
set first;
response=respons5;
trial=5;
keep num response trial groupX colony;
run;
data six;
set first;
response=respons6;
trial=6;
keep num response trial groupX colony;
run;
data new;
set one two three four five six;
run;
(end of data entry)
```

Beginning of regression

```
proc genmod data=new;
class num;
model response= trial trial*groupX / d=b;
repeated subject=num /type=mdep(5) /*may also use
ar(1) or cs */;
run;
```

Note on interpretation of results

SAS gives two sets of regression parameters: use the second set. They are the adjusted estimates from generalized estimating equations (GEE). The first set of parameters is not adjusted for the correlation structure.

References

Ben-Shahar Y, Thompson CK, Hartz SM, Smith BH, Robinson GE (in press) Performance on a reversal learning test is associated with division of labor in the honey bee, *Apis mellifera*. Anim Cogn

Bhagavan S, Benatar S, Cobey S, Smith BH (1994) Effect of genotype but not of age or caste on olfactory learning performance in the honey bee, *Apis mellifera*. Anim Behav 48:1357–1369

Bitterman ME, Menzel R, Fietz A, Schafer S (1983) Classical conditioning of proboscis extension in honeybees (*Apis mellifera*). J Comp Psychol 97:107–119

Dunlop DD (1994) Regression for longitudinal data: a bridge from least squares regression. Am Stat 48:299–303

Gerber B, Geberzahn N, Hellstern F, Klein J, Kowalski O, Wustenberger D, Menzel R (1996) Honey bees transfer olfactory memories established during flower visits to a proboscis extension paradigm in the laboratory. Anim Behav 52:1079–1085

Gerber B, Wustenberg D, Schutz A, Menzel R (1998) Temporal determinants of olfactory long-term retention in honeybee classical conditioning: nonmonotonous effects of the training trial interval. Neurobiol Learn Mem 69:71–78

Littell RC, Miliken GA, Stroup WW, Wolfinger RD (1996) SAS System for Mixed Models. SAS Institute, Cary

Menzel R (1999) Memory dynamics in the honeybee. J Comp Physiol A 185:323–340

Pelz C, Gerber B, Menzel R (1996) Odorant intensity as a determinant for olfactory conditioning in honeybees: roles in discrimination, overshadowing and memory consolidation. J Exp Biol 200:837–847

Ray S, Ferneyhough B (1997) The effects of age on olfactory learning and memory in the honey bee *Apis mellifera*. Neuroreport 8:789–793

Ray S, Ferneyhough B (1999) Behavioral development and olfactory learning in the honeybee (*Apis mellifera*). Dev Psychobiol 34:21–27