



## Average group behavior does not represent individual behavior in classical conditioning of the honeybee

Evren Pamir, Nelay Kumar Chakroborty, Nicola Stollhoff, et al.

*Learn. Mem.* 2011 18: 733-741

Access the most recent version at doi:[10.1101/lm.2232711](https://doi.org/10.1101/lm.2232711)

---

### References

This article cites 29 articles, 4 of which can be accessed free at:  
<http://learnmem.cshlp.org/content/18/11/733.full.html#ref-list-1>

### Email alerting service

Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or [click here](#)

---

---

To subscribe to *Learning & Memory* go to:  
<http://learnmem.cshlp.org/subscriptions>

---

## Research

# Average group behavior does not represent individual behavior in classical conditioning of the honeybee

Evren Pamir,<sup>1,2,4</sup> Nelay Kumar Chakroborty,<sup>3</sup> Nicola Stollhoff,<sup>3</sup> Katrin Barbara Gehring,<sup>3</sup> Victoria Antemann,<sup>3</sup> Laura Morgenstern,<sup>3</sup> Johannes Felsenberg,<sup>3</sup> Dorothea Eisenhardt,<sup>3</sup> Randolph Menzel,<sup>3</sup> and Martin Paul Nawrot<sup>1,2</sup>

<sup>1</sup>Bernstein Center for Computational Neuroscience, Berlin, 10115 Germany; <sup>2</sup>Neuroinformatics & Theoretical Neuroscience, Institute of Biology, Freie Universität Berlin, 14195 Germany; <sup>3</sup>Institute of Biology–Neurobiology, Freie Universität Berlin, 14195 Germany

Conditioned behavior as observed during classical conditioning in a group of identically treated animals provides insights into the physiological process of learning and memory formation. However, several studies in vertebrates found a remarkable difference between the group-average behavioral performance and the behavioral characteristics of individual animals. Here, we analyzed a large number of data (1640 animals) on olfactory conditioning in the honeybee (*Apis mellifera*). The data acquired during absolute and differential classical conditioning differed with respect to the number of conditioning trials, the conditioned odors, the intertrial intervals, and the time of retention tests. We further investigated data in which animals were tested for spontaneous recovery from extinction. In all data sets we found that the gradually increasing group-average learning curve did not adequately represent the behavior of individual animals. Individual behavior was characterized by a rapid and stable acquisition of the conditioned response (CR), as well as by a rapid and stable cessation of the CR following unrewarded stimuli. In addition, we present and evaluate different model hypotheses on how honeybees form associations during classical conditioning by implementing a gradual learning process on the one hand and an all-or-none learning process on the other hand. In summary, our findings advise that individual behavior should be recognized as a meaningful predictor for the internal state of a honeybee—irrespective of the group-average behavioral performance.

Learning and memory formation in vertebrates and invertebrates have been studied on the basis of a large range of classical and operant conditioning paradigms. Typically, the interpretation of experimental results relies on performance measures that were derived by averaging over behavioral observations from identically treated animals. However, several studies have recognized the inadequacy of group-average measures to capture the characteristics of individual behavior and, consequently, the learning-induced changes in individual brains (Krechevsky 1932; Restle 1965; Hanson and Killeen 1981; Estes 2002; Brown and Heathcote 2003; Cousineau et al. 2003). Most notably, Gallistel et al. (2004) found that the gradually increasing learning curve observed in many vertebrate learning paradigms reflected an artifact of group averaging. The behavioral performance of individuals appeared to be characterized by an abrupt and often step-like increase in the level of response.

To our knowledge and in contrast to the vertebrate literature (see Gallistel et al. 2004), surprisingly little is known of a possibly heterogeneous expression of behavior for the most frequently applied invertebrate conditioning paradigms. For the fruit fly (*Drosophila melanogaster*) it appears to be common sense that the group-average behavioral measures adequately represent the probabilistic expression of behavior in individuals—a notion that goes back to an early study by Quinn et al. (1974).

In the following, we focus on a classical conditioning paradigm in the honeybee (*Apis mellifera*)—the reward-based olfactory conditioning of the proboscis extension response (Takeda 1961; Bitterman et al. 1983). In this paradigm, a group of harnessed

honeybees is individually trained by forward pairings of an olfactory stimulus (CS) with a sucrose reward (US). The conditioned response (CR), namely, the extension or nonextension of the proboscis at each of these conditioning trials, is typically documented in a binary form. Additional characteristics of the proboscis extension can be captured by using a high-speed video or by recording the electromyogram of the muscles involved in the proboscis extension (Rehder 1987; Smith and Menzel 1989a,b). Importantly, conditioning of the proboscis extension response allows one to simultaneously measure brain activity by means of electrophysiology or calcium imaging (Giurfa 2007; Menzel et al. 2007), giving access to the neuronal correlates of learning and memory. The molecular mechanisms of memory consolidation in the honeybee brain can be studied by combining classical conditioning with *in vivo* pharmacological interventions or post-mortem biochemical analysis (Stollhoff et al. 2005; Eisenhardt 2006; Schwärzel and Müller 2006).

In the present study, we determine the behavioral characteristics of individual honeybees during classical conditioning on the basis of a large collection of experimental data. Specifically, we ask: (1) How well does the group-average behavioral performance represent the behavioral performance of individuals during absolute conditioning, differential conditioning, and extinction? (2) How do individual animals learn during absolute conditioning, and how does this learning translate into behavior? The Results section is divided into two parts. In the first part, we answer the first question by means of an exploratory data analysis. In the second part, we consider three different answers to the second question by evaluating the eligibility of three different generative models. The implications of our findings for the analysis and interpretation of behavioral and physiological data from the honeybee are discussed.

#### <sup>4</sup>Corresponding author.

E-mail [evren.pamir@bccn-berlin.de](mailto:evren.pamir@bccn-berlin.de).

Article is online at <http://www.learnmem.org/cgi/doi/10.1101/lm.2232711>.

## Results

### Experimental data

We analyzed data from a number of independent experiments that were conducted at the Institute of Biology–Neurobiology of the Freie Universität Berlin between the years 1999 and 2009. None of the experiments were originally designed for the purpose of the present study. We collected a total of 17 data sets comprising 1640 animals (see Table 1). The animals in data sets 1–15 were trained by using an absolute classical conditioning protocol, while the animals from the data sets 16 and 17 were differentially conditioned (see Materials and Methods). The data sets differ with respect to the number of conditioning trials  $m$ , the temporal inter-trial interval during acquisition (ITI), the time-point of the memory retention test ( $T$ ), and the odor that was used as the conditioned stimulus (CS). A subset of animals ( $n = 217$ ) from data set 1 was subjected to an extinction protocol.

### For absolute conditioning, the group CR probability does not represent individual behavior

For each of the data sets 1–15 (absolute classical conditioning), we asked how well the group CR probability  $P(x_t = 1)$  in a given trial  $t$  represented the CR probability of subgroups of animals. We compared the behavioral performances of two disjoint subgroups defined by their response in the previous trial ( $x_{t-1} = 1$  or  $x_{t-1} = 0$ ). Figure 1A shows an example of the standard data analysis (see

**Table 1.** Data analyzed from absolute (data sets 1–15) and differential (data sets 16–17) conditioning experiments

Data set	$m$	ITI (min)	$T$ (h)	$N$	Odor	Data origin
1	3	10	24	517	Clove oil	Stollhoff et al. 2005
2	3	10	25	98	Clove oil	Stollhoff et al. 2005
3	3	10	26	113	Clove oil	Stollhoff et al. 2005
4	3	10	28	92	Clove oil	Stollhoff et al. 2005
5	3	10	48	85	Clove oil	Stollhoff et al. 2005
6	3	10	72	94	Clove oil	Stollhoff et al. 2005
7	3	10	24	87	Clove oil	KB Gehring and D Eisenhardt, unpubl.
8	3	2	24	58	Clove oil	L Morgenstern, J Felsenberg, and D Eisenhardt, unpubl.
9	3	2	24, 48	64	Clove oil	V Antemann, J Felsenberg, and D Eisenhardt, unpubl.
10	4	30	1, 24	56	Isoamyl Acetate	NK Chakroborty, unpubl.
11	4	30	1, 24	66	Isoamyl Acetate	NK Chakroborty, unpubl.
12	5	30	1, 24	48	7-pentadecene	NK Chakroborty, unpubl.
13	5	30	1, 24	37	6-pentadecene	NK Chakroborty, unpubl.
14	12	0,5	96	63	Hexanol	Menzel et al. 2001
15	12	15	96	64	Hexanol	Menzel et al. 2001
16	6,6'	15	1, 24	48	1-Hexanal, 1-octanol'	NK Chakroborty, unpubl.
17	6,6'	15	1, 24	50	1-Octanol, 1-hexanal'	NK Chakroborty, unpubl.

The 17 data sets differed with respect to several experimental parameters. ( $m$ ) Number of conditioning trials in the acquisition session; ITI intertribal interval in the acquisition session; ( $T$ ) time of the retention test; ( $N$ ) number of animals in each data set. For the data sets 16 and 17, the apostrophe indicates CS– trials. The data were contributed by the individual co-investigators as indicated.

Materials and Methods) for the data set 10. While performing this analysis for data sets 1–15 we found a remarkably uniform pattern during conditioning and in the memory retention test (Fig. 1B,C): Animals that responded in the previous trial  $t - 1$  always had a higher chance for responding in trial  $t$  than animals that did not respond in trial  $t - 1$ . The differences between the subgroup CR probabilities  $P(x_t = 1|x_{t-1} = 1)$  and  $P(x_t = 1|x_{t-1} = 0)$  were statistically significant ( $\chi^2$  test,  $\alpha = 0.05$ ) in 39 out of 39 cases during conditioning (Fig. 2B), and in 15 out of 20 cases during memory retention (Fig. 2C). For each of the 15 data sets we also computed the mean probabilities for the observations ( $x_t = 1|x_{t-1} = 0$ ) and ( $x_t = 1|x_{t-1} = 1$ ) across all possible trials. Computing the average of these mean probabilities across data sets 1–15 yielded ( $\pm$ SD):  $P(x_t = 1|x_{t-1} = 1) = 0.89 \pm 0.05$  and  $P(x_t = 1|x_{t-1} = 0) = 0.43 \pm 0.14$ .

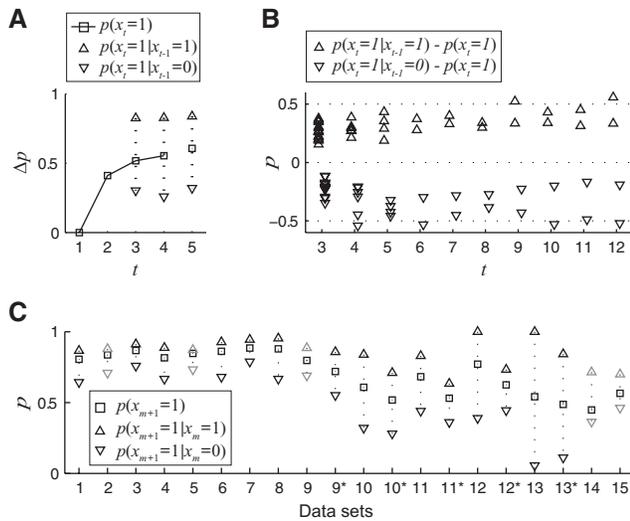
### Differential conditioning is consistent with absolute conditioning

We next asked whether the results obtained for absolute conditioning also applied to data from differential conditioning experiments. For the data sets 16 and 17 we separately analyzed the CS+ and CS– trials using the standard data analysis (see Materials and Methods). We found that animals that responded to the CS+ (CS–) in the previous trial  $t - 1$  ( $t' - 1$ ) had a higher chance for responding to the CS+ (CS–) in trial  $t$  ( $t'$ ) than animals that did not respond to the CS+ (CS–) in the previous trial (Fig. 2). Once animals responded to the CS+, or once animals did not respond to the CS– anymore, the response probability remained rather constant in the following trials. Thus, the group CR probability does not adequately represent the behavioral characteristics of individual animals during differential conditioning.

### The cessation of the CR is abrupt in individuals during extinction

A subset of animals ( $n = 217$ ) from data set 1 was presented in five extinction trials 24 h after conditioning (Fig. 3A,B). After the extinction session, the animals were divided into five groups, each of which was tested for spontaneous recovery from extinction at a different time point (Fig. 3C; see also Fig. 3 in the original study by Stollhoff et al. 2005). By using the standard data analysis (see Materials and Methods) we asked whether the heterogeneous expression of behavior observed during conditioning was also present during extinction and subsequent memory retention. For all trials, we again found that animals that responded in the previous trial  $t - 1$  had a higher chance for responding in the trial  $t$  than animals that did not respond in the previous trial (Fig. 3A–C). The analysis also revealed that once individual animals had ceased to respond during extinction ( $t = 4$ –8), they had a high chance of remaining nonresponding in subsequent extinction trials (Fig. 3B). Hence, analogous to our findings for the acquisition of a CR during conditioning, the cessation of the CR during extinction is characterized by an abrupt rather than gradual change of CR probability in individuals.

On the basis of our previous analysis, we hypothesized that the presence or absence of the CR in the retention test ( $t = m + 1$ ) may serve as a binary indicator for the actual learning success of individual animals. Consequently, showing a CR on this trial ( $t = 4$ ) may indicate that an associative memory has been induced by the conditioning procedure, while not showing a CR on this trial may indicate that no association or only a poor association has been formed. To test our hypothesis we divided the 217 animals into two disjoint subgroups that were defined by their response on trial  $t = 4$  ( $x_4 = 1$  and  $x_4 = 0$ ). As expected, our heuristic selection criterion yielded a steady performance difference



**Figure 1.** For data sets 1–15, the group CR probability is a poor estimate for individual behavior during conditioning and in the memory-retention tests. Animals that responded in the previous trial  $t - 1$  always showed a higher probability for responding in trial  $t$  than animals that did not respond in the previous trial  $t - 1$ . (A) Group CR probabilities  $P(x_t = 1)$  and the two subgroup CR probabilities  $P(x_t = 1|x_{t-1} = 1)$  and  $P(x_t = 1|x_{t-1} = 0)$  for data set 11. The data were analyzed by the standard data analysis (see Materials and Methods). Black triangles indicate a significant difference between the two subgroup probabilities in the respective trial ( $\alpha = 0.05$ ). Gray triangles indicate a nonsignificant difference. (B) Subgroup CR probabilities  $P(x_t = 1|x_{t-1} = 1)$  and  $P(x_t = 1|x_{t-1} = 0)$  computed for all conditioning trials and for all 15 data sets. For display, we subtracted the group CR probabilities from the subgroup CR probabilities. (C) Subgroup CR probabilities during memory retention conditioned on the outcome of the final conditioning trial computed for data sets 1–15. For data sets 9–13, the asterisk indicates the second memory-retention test (see Table 1).

between the two subgroups during conditioning (Fig. 3D): the animals in the first subgroup ( $x_4 = 1$ ) had a higher response probability than the animals in the second subgroup ( $x_4 = 0$ ). This performance difference also persisted during extinction and memory retention (Fig. 3E,F), which may provide evidence for the heterogeneity of the 217 conditioned animals with respect to the associative memories formed. We tested the differences of the two subgroups for statistical significance by a  $\chi^2$  test and found that the differences were always significant during acquisition ( $t = 2, 3$ ) and extinction ( $t = 5-8$ ), but only significant once in the subgroups tested for memory retention (25 h). It should be noted that the test power was much lower in the retention tests because the group sizes were reduced by a factor of 4.

### The simple learning-curve model must be rejected

Given the typical asymptotic rise of the group CR probability in a classical conditioning experiment, one may assume that the shape of the curve reflects a gradual increase of associative strength  $AS(t)$  across trials in individuals. One may further assume that all animals are identical with respect to this learning process and that individual animals express a CR at trial  $t$  with probability  $P^{LCM1}(x_t = 1) = AS(t)$ . We tested the eligibility of this model hypothesis, termed the simple learning-curve model, on data sets 1–15 (see Fig. 4A and Materials and Methods for the full model description). While this model fit the group CR probabilities very accurately (Fig. 6A, below), it did not capture the CR probabilities at the level of subgroups. As we expected from our data analysis in the previous sections, the empirical subgroup

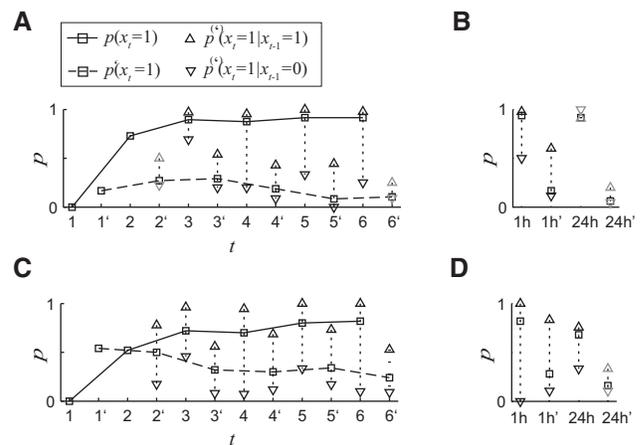
probabilities  $P(x_t = 1|x_{t-1} = 1)$  and  $P(x_t = 1|x_{t-1} = 0)$  significantly deviated (two-tailed binomial test,  $\alpha = 0.05$ ) from the model predictions in 89 out of 104 cases (Fig. 5A).

### The learning-curve model can be saved by two additional features

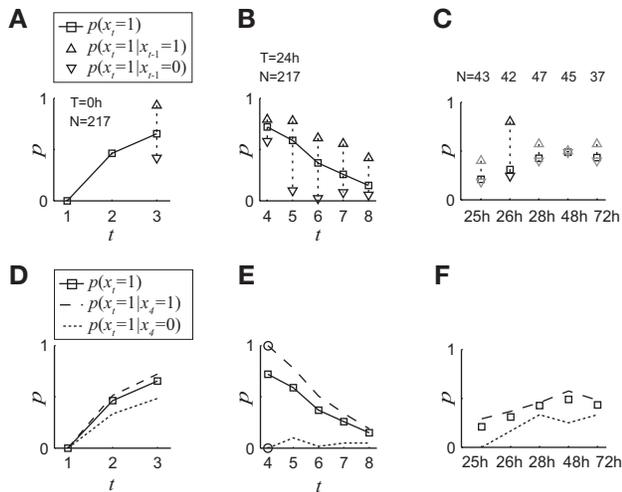
In this section we maintain the previous notion of a gradual increase of associative strength in individuals across conditioning trials; however, we add the possibility that the rate of learning may not be the same in all individuals. Furthermore, we use a simple performance rule, effectively a threshold criterion, which maps the associative strength in individuals to the CR probability (see Fig. 4B and Materials and Methods for a full description of this model). By fitting this extended learning-curve model to data sets 1–15 we found that it provided a possible explanation for the observed behavioral sequences. The empirical subgroup probabilities  $P(x_t = 1|x_{t-1} = 1)$  and  $P(x_t = 1|x_{t-1} = 0)$  deviated significantly from the model predictions in only nine out of 104 cases (Fig. 5B). We also noticed that the estimated probability distributions were skewed to the highest learning-rate interval (see Fig. 7, below). Computing an average probability distribution over data sets 1–9, we found that the highest resolvable learning-rate interval is adopted by 48% of the animals, while the second, third, and fourth highest intervals are adopted by 26%, 18%, and 8% of the animals, respectively. For data sets 10–15, in which the lower learning-rate intervals were resolvable, we always found another peak of the distribution at the respective lowest interval.

### A two-state hidden Markov model captures the behavioral sequences

Our data analysis revealed that once animals were responding in a given trial, they had a very high probability of again responding in



**Figure 2.** During differential conditioning, animals that responded to the CS+ (CS−) in the previous trial  $t - 1$  ( $t' - 1$ ) showed a higher probability for responding to the CS+ (CS−) in trial  $t$  ( $t'$ ) than animals that did not respond to the CS+ (CS−) in the previous trial. The apostrophe indicates CS− trials. (A) Group and subgroup probabilities during conditioning for data set 16. Behavioral responses in CS+ and CS− trials were independently analyzed by the standard data analysis (see Materials and Methods). Black triangles indicate a significant difference between subgroup probabilities ( $\alpha = 0.05$ ). Gray triangles indicate a nonsignificant difference. (B) Animals from data set 16 were repeatedly tested for memory retention by presenting the CS+ and the CS− at  $T = 1$  h and 24 h. The subgroup probabilities were conditioned on the response to the last CS+ (CS−) presentation during conditioning. (C,D) Same analysis as in A and B, but for a different group of animals (data set 17), which was differentially conditioned to the reversed odor-reward contingencies.



**Figure 3.** A subset of animals ( $n = 217$ ) from data set 1 was presented to five extinction trials 24 h after conditioning (A,B). After extinction ( $t > 8$ ), animals were divided into five groups, each of which was tested for spontaneous recovery at a different time point (C). (A–C) During acquisition, extinction, and the retention tests, animals that responded in the previous trial  $t - 1$  showed a higher probability for responding in trial  $t$  than animals that did not respond in the previous trial. The standard data analysis was used without change from trial  $t = 1 \dots 8$  (see Material and Methods). The subgroup probabilities in the five retention tests were conditioned on the response in trial  $t = 8$ . Black triangles indicate a significant difference between subgroup probabilities ( $\alpha = 0.05$ ). Gray triangles indicate a nonsignificant difference. (D–F) Group CR probabilities as in A, B, and C. The two dashed curves show the response probabilities for the two disjoint subgroups defined by their behavior in trial  $t = 4$ . The employed selection criterion (black circles) results in a steady performance difference between the two subgroups of animals in all other trials.

the next trial. We hypothesized that this type of response behavior could be adequately captured by a hidden Markov model with two hidden states,  $k = 1$  and  $k = 2$ , and two possible observations,  $x_t = 0$  and  $x_t = 1$  (see Fig. 4C and Materials and Methods for a full description of this model). Fitting the two-state hidden Markov model to data sets 1–15, we found that the parameter estimation yielded highly similar results for each of the data sets. Furthermore, for each data set the estimated parameters did not depend on the chosen starting values. We report here their mean values averaged over data sets 1–15 ( $\pm$  SD):  $P(k_1 = 1) = 1.00$ ,  $P(k_t = 2|k_{t-1} = 1) = 0.43 \pm 0.16$ ,  $P(k_t = 2|k_{t-1} = 2) = 0.96 \pm 0.04$ ,  $P(x_t = 0|k_t = 1) = 1.00 \pm 0.01$ ,  $P(x_t = 1|k_t = 2) = 0.94 \pm 0.04$ . These values express the following prototypic scenario: At  $t = 1$  all animals are in the state  $k = 1$ . Animals in this state have a very low probability for showing a CR and a moderate probability for making a transition into the state  $k = 2$ . The second state is characterized by a very high probability for showing a CR and a very high probability for remaining in this state. As a consequence, once an individual animal has made a transition to the second state  $k = 2$ , it will show a stable expression of the CR in the subsequent trials. As in the case of the extended learning-curve model, we found that the hidden Markov model provided a possible explanation for the observed behavioral sequences. The empirical subgroup probabilities  $P(x_t = 1|x_{t-1} = 1)$  and  $P(x_t = 1|x_{t-1} = 0)$  only significantly deviated from the model predictions in 11 out of 104 cases (Fig. 5C).

### Model comparison by cross-validation

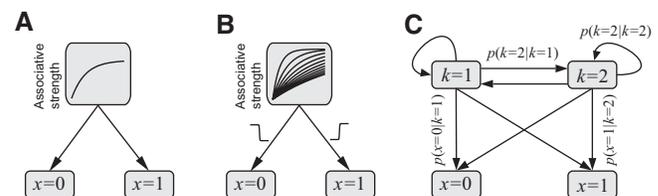
The more complex an explanation, the better it can capture a given set of data. Simple explanations, however, are often the

best. In terms of complexity, the simple learning-curve model is the simplest, comprising only two free parameters. In turn, the extended learning-curve model is the most complex because it uses an unconstrained empirical probability estimate for the learning-rate distribution. For a data set with  $L$  trials, the extended learning model has  $L + 1$  free parameters. With four free parameters, the hidden Markov model lies between the other two models in terms of complexity. By fitting the three models to data sets 1–15 we found that the extended learning-curve model and the hidden Markov model could capture the behavioral observations, while the simple learning-curve model could not (Figs. 5, 6). However, this finding may result from their higher complexity rather than from their higher explanatory power. To directly compare the eligibility of the three models, we computed their log likelihoods by a cross-validation algorithm for each of the data sets 1–15 (see Materials and Methods). This procedure confirmed that the simple learning-curve model had to be rejected, since it performed worse than each of the other two models in all cases (see Table 2). We also found that the extended learning-curve model was more likely than the hidden Markov model in 12 out of 15 cases.

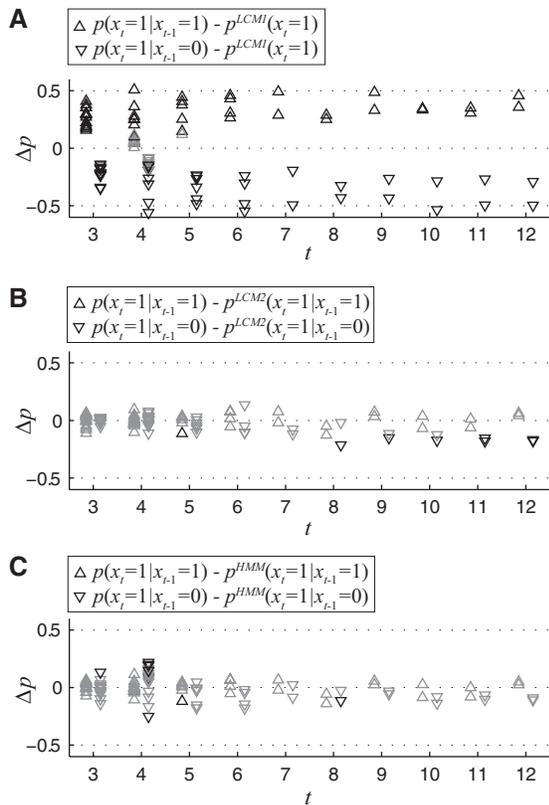
### Discussion

#### The group CR probability is a poor estimate for the behavior of individual honeybees

Gallistel et al. (2004) analyzed several vertebrate learning paradigms for how well the group-average performance measures represented the behavioral performance of individuals. They found that the gradual and negatively accelerated increase seen at the population level defining the empirical learning curve was an artifact of group averaging. Individual behavior was characterized by an abrupt and often step-like change in the level of responding. Since at each time point the population comprises two types of animals, those that have acquired the CR, and those that have not, the investigators suggest describing the experimental data by individual onset latencies rather than by group-average measures.



**Figure 4.** Three model hypotheses for the observed binary sequences of conditioned responses. (A) The simple learning-curve model (LCM1) hypothesizes that the gradually increasing group CR probability  $P(x_t = 1)$  reflects the gradually increasing associative strength in individual animals. All animals are the same and at a given trial  $t$  express the CR with probability  $P^{\text{LCM1}}(x_t = 1)$ . (B) The extended learning-curve model (LCM2) assumes that individual animals in a given group can differ with respect to their learning rates, which is described by a discrete probability distribution  $P(\varepsilon_i)$ . The probability for observing or not observing a CR is described by a simple performance function that is the same for all animals of the group. (C) The hidden Markov model assumes that individual animals can be in one of two discrete hidden states. At the beginning of the conditioning experiment all animals are in the naive or unlearned state  $k = 1$ . At each conditioning trial, animals in this state have a certain success rate  $P(k = 2|k = 1)$  for making a transition to the learned state  $k = 2$ . The probability  $P(k = 2|k = 2)$  for remaining in the learned state is typically very high. The probability  $P(x = 1|k = 2)$  for expressing a CR is high in the state  $k = 2$ , while the CR probability  $P(x = 1|k = 1)$  is low in the state  $k = 1$ . (See Materials and Methods for a description of how parameters are estimated for the three models.)



**Figure 5.** Differences between the empirical subgroup's CR probabilities and the respective model estimations for data sets 1–15. The differences were tested for statistical significance with a two-tailed binomial test. Black triangles denote statistical significance ( $\alpha = 0.05$ ). Gray triangles denote no statistical significance. (A) The simple learning-curve model (LCM1) cannot capture the experimental data at the level of subgroups. (B,C) Both the extended learning-curve model (LCM2, B) and the hidden Markov model (HMM, C) provide a good fit for the experimental data.

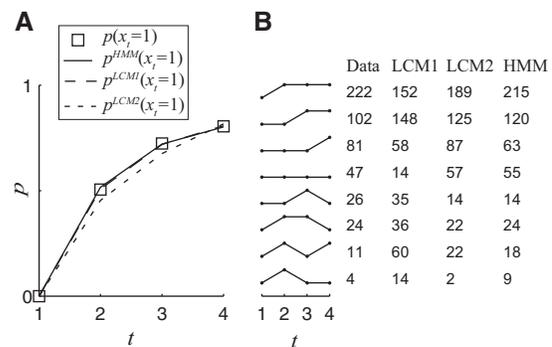
As we have shown, the case is very similar for classical conditioning of the proboscis extension response in the honeybee. For both absolute and differential conditioning, we found a very high probability that a bee would again extend its proboscis, given that it had done so in the previous trial. This high probability clearly excludes the adequacy of using the gradually increasing group CR probability to represent the response probability of individuals during the training phase. Most importantly, we found that the heterogeneous expression of behavior observed during the conditioning phase persisted during memory retention, which indicates a heterogeneity in a group of identically treated animals with respect to long-term memory formation. Analyzing data in which animals were tested for spontaneous recovery from extinction provided additional evidence for this notion. Interestingly, we also observed a rapid and stable change of the response probabilities when honeybees had to learn that stimuli were not followed by a reward, as was the case for the unrewarded stimuli presentations during differential conditioning and extinction. Thus, once individual animals did not respond to an unrewarded stimulus in a given trial, they had a high probability of not responding in subsequent trials. In summary, individual behavior is characterized by abrupt and stable changes in response probabilities, contrary to the gradual changes in group CR probability observed at the population level.

## How representative is the mean in the fruit fly?

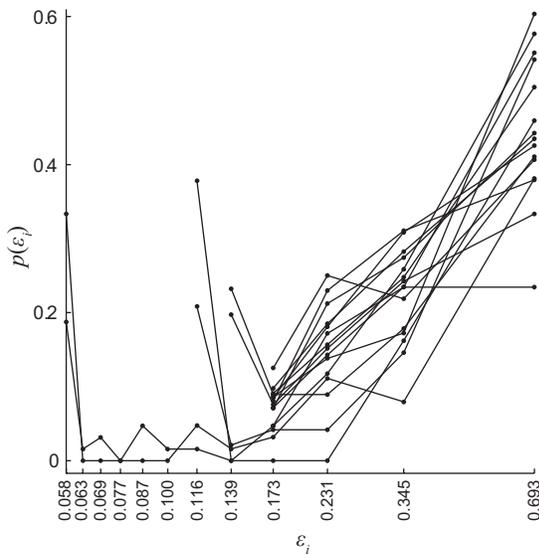
Our results differ from the commonly held notion of a homogeneous expression of conditioned behavior in the fruit fly, as first reported in a study by Quinn et al. (1974). For a binary choice olfactory-conditioning paradigm, Quinn et al. (1974) asked whether the group-average performance observed after conditioning arose from some heterogeneity in the population, or whether it was due to a stochastic component in the behavior of all of the flies. To answer this question, the investigators separated the flies that made the correct choice from those that did not, and 24 h later retrained and retested each group. Since the performance of both groups was the same, the investigators concluded that the expression of behavior was probabilistic in each fly, and that there was no evidence for an intelligent subset of the population. A recent study by Chabaud et al. (2010) made a more ambiguous observation when tracking the choice behavior of individual fruit flies during memory retention. For two types of training protocols, the investigators found that the expression of behavior was probabilistic during the test, because the choice of individual flies at the end of the test was not determined by their first choice in the test. However, for a third training protocol, the final memory score was determined by the first choice, yielding a bipolar distribution of the individual memory scores. It remains to be more thoroughly investigated whether serial correlations in the behavior of individuals during training, and between training and testing as reported here, is a specific result for the honeybee or whether it also applies to other invertebrates.

## How do individual honeybees learn?

For the honeybee, several experimental parameters and conditions have been described as having a decisive effect on the strength, specificity, and stability of associative memories by comparing group-average CR probabilities in differently treated groups (for the number of conditioning trials and intertrial intervals, see Bitterman et al. 1983; Sandoz et al. 1995; Gerber et al. 1998; Menzel 1999; Menzel et al. 2001; for bee age or season, see Behrends and Scheiner 2010; Hadar and Menzel 2010; for US strength, see Menzel et al. 2001; Scheiner et al. 2004). To pick



**Figure 6.** Illustration of the performance differences between the simple learning-curve model (LCM1), the extended learning-curve model (LCM2), and the hidden Markov model (HMM) for data set 1 ( $n = 517$ ). (A) All three models can describe the behavioral data at the population level. (B) The data set 1 contains eight different binary behavioral sequences. (Left column) Absolute sequence frequencies as counted in the data. (Middle and right column) Absolute sequence frequencies as computed by the three models rounded to integers. The distribution of sequences in the experimental data is best captured by the hidden Markov model (75 errors), followed by the extended learning-curve model (99 errors) and the simple learning-curve model (252 errors). Errors are the sum over the absolute differences between frequencies in the data and in the respective model prediction.



**Figure 7.** Empirical probability distribution estimates  $P(\varepsilon_i)$  for data sets 1–15. The  $x$ -axis has been scaled logarithmically. The values of  $\varepsilon_i$  on the  $x$ -axis always denote the lower limits of the respective learning-rate intervals (see Materials and Methods, “The extended learning-curve model”).

the most basic experimental parameter, if one group A is trained with more conditioning trials than another group B, this typically yields a higher group CR probability in group A at the end of the conditioning phase, as well as in the memory-retention test compared with group B. The classical interpretation of this canonical observation is that more training yields stronger memories, implicitly assuming that each of the two groups is rather homogeneous with respect to learning and memory formation under given experimental conditions. However, both of the extended learning-curve model and the hidden Markov model are at odds with this line of reasoning.

According to the extended learning-curve model, the population is heterogeneous with respect to learning rates (see Fig. 7). Consequently, in any trial  $t$ , the population will comprise animals with associative strengths below threshold that do not show a CR, and animals with associative strengths above threshold that show a CR with high probability. Hence, the group CR probability reflects the ratio between the sizes of the two subsets of animals at any trial  $t$ . If group A is trained with more trials than group B, then this results in a higher group CR probability at the end of the conditioning phase as well as in the memory retention test in group A, however, not because stronger memories have been induced in individuals of group A, but simply because the ratio between the sizes of the two subsets has been shifted to a larger value. To be more specific, the extended learning-curve model assumes that animals with an associative strength above threshold increase their associative strength with further conditioning trials, but this process does not contribute to the shift of the observed ratio, nor is it visible in the individual behavioral sequences. To conclude that associations in individuals have been strengthened, would at least require nonbinary behavioral measures. For example, when recording the potentials of the M17 muscle involved in the proboscis extension response, Smith and Menzel (1989a,b) observed a gradual rise of the muscle potentials and a gradual decrease of the response latencies with the number of conditioning trials.

The hidden Markov model assumes all-or-none learning, which is that at each trial  $t$ , a given animal has a certain success rate  $P(k=2|k=1)$  for learning a stimulus-reward association in

an all-or-none fashion. Whenever an association has been learned by an animal, this association is highly stable across the remaining trials. According to this view, more training trials increase the number of success opportunities for each animal and, consequently, more training trials will increase the number of animals in the second state. At each trial  $t$ , the group CR probability then simply reflects the ratio between the number of animals in the naive state ( $k=1$ ) and in the learned state ( $k=2$ ); however, the group CR probability does not contain any information about the state of an individual animal. (There are two exceptions: when all animals occupy either the first or the second state). It should be noted that in the hidden Markov model view, individual animals only need one successful trial to establish a stable association. Some experimental support for this possibility comes from a study that showed honeybees can form a stable and long-lasting memory even after a single conditioning trial (Sandoz et al. 1995). One can further hypothesize that the success rate  $P(k=2|k=1)$  depends on a large number of experimental conditions, as well as on intrinsic conditions of the respective animal such as hunger, motivation, hormonal status, health, or age, while the probability for maintaining a formed association is possibly invariant to these factors.

To summarize, the extended learning-curve model and the hidden Markov model represent two alternative hypotheses for the dynamics of associative learning in the honeybee during classical conditioning. Further experiments to test the various consequences of the two hypotheses are underway.

## Implications for data analysis

Our findings advise that individual behavior should be recognized as a meaningful predictor for the internal state of a honeybee—irrespective of the group CR probability. In particular, this suggests that the analysis of parallel behavioral and physiological recordings should be carried out at the level of individual animals. Several studies have demonstrated how this can lead to a more informative analysis. For a rule-learning task in rats, Durstewitz et al. (2010) found that neuronal activity recorded from the prefrontal cortex was in tight temporal relation to behavioral performance shifts in individuals. For the honeybee, two recent studies divided groups of identically treated animals into subgroups of

**Table 2.** Differences between the mean log-likelihoods of the three models after 50 rounds of fourfold cross-validation

Data set	$\log(P^{LCM2}) - \log(P^{LCM1})$	$\log(P^{LCM2}) - \log(P^{HMM})$
1	19.3	0.0
2	3.6	0.4
3	5.4	1.2
4	3.3	0.7
5	2.7	0.5
6	5.0	0.9
7	1.7	0.9
8	4.0	0.8
9	2.3	0.9
10	6.0	-0.2
11	6.3	0.3
12	12.6	1.9
13	14.6	3.4
14	14.0	-20.9
15	31.4	-17.8

(LCM1) Simple learning-curve model; (LCM2) extended learning-curve model; (HMM) two-state hidden Markov model. The extended learning-curve model was chosen as a reference point because it was the best model for 11 of the 15 data sets. The extended learning-curve model performs better than the simple one in all 15 cases, and it also performs better than the hidden Markov model in 11 out of 15 cases.

so-called learners and nonlearners on the basis of a heuristic behavioral selection criterion (Roussel et al. 2010; Rath et al. 2011). Both studies then found differences in the simultaneously recorded CA-imaging signals between the two subgroups. Here, we also suggest that behavioral studies that use *in vivo* pharmacological interventions or post-mortem biochemical analysis of honeybee brains should take into account individual behavior as a possibly discriminative factor.

Our analysis offers several possibilities to illustrate the heterogeneity in a given set of behavioral data.

Adopting the extended learning-curve model, the distribution of learning rates across individuals can be empirically estimated and visualized. Adopting the hidden Markov model, the transition and observation probabilities for the naive and the conditioned state can be determined. Under a model-free perspective, conditional probabilities can be computed as demonstrated in the exploratory data analysis part of this study. Commented Matlab code for all analysis carried out here is available on request.

## Materials and Methods

### Absolute classical conditioning

The exact experimental procedure used during conditioning of the proboscis extension response has been described elsewhere (Menzel et al. 2001; Stollhoff et al. 2005; Felsenberg et al. 2011). During absolute classical conditioning (data sets 1–15), honeybees were exposed to  $m$  forward pairings of the conditioned stimulus (CS, odor) with the unconditioned stimulus (US, sucrose). The CS and US durations, as well as the CS–US overlaps slightly differed in the data sets (data sets 1–13: 5 sec CS, 4 sec US, 2 sec CS–US overlap; data sets 14–15: 4 sec CS, 3 sec US, 1 sec CS–US overlap). Since the CS onset preceded the US onset by a few seconds, the occurrence of the proboscis extension during this time span was documented in a binary form as the CR. For a certain trial  $t$ , we denote the presence (absence) of the CR with  $x_t = 1$  ( $x_t = 0$ ). After conditioning, memory retention was tested by exposing the bees to the CS alone at time  $T$ . The animals from data sets 9–13 were tested two times at  $T = 1$  h and at  $T = 24$  h. In this study we only included animals that (1) did not respond to the first CS during acquisition, (2) survived the entire experiment, and (3) showed the proboscis extension response elicited by sucrose feeding at the very end of the experiment. We made one exception to this rule: data sets 14 and 15 comprise 63 and 64 animals that were conditioned, but only 29 and 23 animals were tested at  $T = 94$  h.

### Differential conditioning

In the differential conditioning paradigm (data sets 16 and 17) honeybees experienced two different odors during conditioning, the first one being rewarded (CS+) and the second one being unrewarded (CS–). Each data set comprised 12 conditioning trials (6 CS+ and 6 CS– in alternating order, starting with CS+, ITI = 14 min), as well as a retention test for both odors at  $T = 1$  h and  $T = 24$  h. In the data set 16, 1-hexanal and 1-octanol were used as CS+ and CS–, respectively, while in the data set 17, the odor-reward contingencies were reversed. The CS+ and CS– durations were 5 sec, the US duration was 4 sec, and the overlap between the CS+ and the US was 2 sec.

### Standard data analysis

Each data set was independently analyzed by the following standard procedure: In a given trial  $t$  we distinguished between group CR probabilities and subgroup CR probabilities. To compute the group CR probability  $P(x_t = 1)$  we divided the number of bees that showed a CR in trial  $t$  by the total number of bees  $N$ . Plotting  $P(x_t = 1)$  against trial order  $t$  results in the so-called learning curve (see, e.g., Figs. 1A, 3A). The subgroup probability  $P(x_t = 1|x_{t-1} = 1)$  is conditioned on the expression of the CR in trial  $t - 1$ . It was computed by dividing the number of bees that showed a

CR in both trials  $t$  and  $t - 1$  by the number of bees that showed a CR in trial  $t - 1$ . The subgroup probability  $P(x_t = 1|x_{t-1} = 0)$  was computed by dividing the number of bees that showed a CR in trial  $t$  but not in trial  $t - 1$  by the number of bees that did not show a CR in trial  $t - 1$ . The difference between the two subgroup CR probabilities was tested for statistical significance by a  $\chi^2$  test. The significance level was set to 0.05.

We specify three amendments to this analysis: (1) For the experiments in which animals were presented to two retention tests (data sets 9–13, 16, 17), the subgroup CR probabilities in the tests were always conditioned on the response in the last conditioning trial ( $t = m$ ). (2) For the data sets 16 and 17 (differential conditioning), the analysis was separately used for CS+ and CS– trials. We introduced an apostrophe in the respective terms to indicate CS– trials. (3) For the data from the extinction experiment (a subset of 217 animals from data set 1) the analysis was used for trials  $t = 1$ –8 without any changes. The subgroup probabilities in the five retention tests were conditioned on the respective behavior in trial  $t = 8$ .

### The simple learning-curve model

The simple learning-curve model (LCM1) hypothesizes that the gradual rise of the CR probability  $P(x_t = 1)$  observed at the population level reflects the gradual rise of associative strength in individuals (see Fig. 4A). Furthermore, all animals in a given population of identically treated animals are assumed to be identical with respect to learning and behavioral performance probabilities. The gradual increase of associative strength (AS) across trials is defined by the learning-rate  $\varepsilon$  and the asymptotic value  $r$ . For each data set the two parameters  $\varepsilon$  and  $r$  are computed by fitting the equation

$$AS(t) = r(1 - \exp(-\varepsilon(t - 1))) \quad (1)$$

to the group CR probabilities  $P(x_t = 1)$  on the range  $t = 1 \dots m + 1$  by nonlinear regression. For the data sets 9–13, in which animals were tested twice, the trial  $t = m + 1$  equaled the first test, while the second test was discarded. (It should be noted that for estimating the parameters of the other two models [see below] we also did not make use of the second test.) Finally, the expression of behavior in each animal is assumed to be probabilistic, with probability  $P^{LCM1}(x_t = 1) = AS(t)$  for expressing a CR, and probability  $P^{LCM1}(x_t = 0) = 1 - AS(t)$  for not expressing a CR, respectively. It should be noted that the equation for the associative strength used here has an identical outcome as the Rescorla-Wagner delta rule for associative learning (Rescorla and Wagner 1972).

### The extended learning-curve model

The extended learning-curve model (LCM2) extends the simple learning-curve model by the following two features (see Fig. 4B): First, learning rates can differ within a population of identically treated animals, which is described by a discrete probability distribution  $P(\varepsilon_i)$ . Second, for a given animal with learning-rate  $\varepsilon_i$  the probability for expressing or not expressing a CR at a given trial  $t$  is defined by a simple threshold function:

$$P^{LCM2}(x_t = 1) = \begin{cases} 0 & \text{for } AS(t) < 0.5 \\ K & \text{for } AS(t) \geq 0.5 \end{cases} \quad \text{and} \quad (2)$$

$$P^{LCM2}(x_t = 0) = 1 - P^{LCM2}(x_t = 1),$$

with

$$AS(t) = 1 - \exp(-\varepsilon_i(t - 1)). \quad (3)$$

Setting the threshold to 0.5 and the asymptotic value  $r$  to unity ensures an optimal dynamic range for any possible distribution of learning rates. According to the extended learning-curve model, different animals can have different associative strengths at a given trial  $t$ ; however, the behavioral performance rule (Eq.

2) is assumed to be the same for all animals: Animals with an associative  $AS(t) < 0.5$  will not show a CR at trial  $t$ , while animals with  $AS(t) \geq 0.5$  will show a CR with probability  $K$ . The parameter  $K$  is a heuristic estimate and equals the mean probability for making the observation ( $x_t = 1 | x_{t-1} = 1$ ) across all possible trials and animals in a data set. Given these model specifications, it follows that the probability distribution of learning rates has to be estimated on the discrete learning-rate intervals:

$[-\ln(1-0.5)/L, -\ln(1-0.5)/(L-1)], [-\ln(1-0.5)/(L-1), -\ln(1-0.5)/(L-2)], \dots, [-\ln(1-0.5)/2, -\ln(1-0.5)/1], [-\ln(1-0.5)/1, \infty]$ , where  $L$  is the maximum number of trials in a data set. (For example, for a data set with  $m$  conditioning trials and one test trial  $L$  equals  $m+1$ .) Any two animals with learning rates from the same interval cannot be dissociated because they will have equal probabilities for any binary behavioral sequence of length  $L$ . The highest learning-rate interval contains all possible learning rates that would result in an associative strength equal to or larger than the threshold of 0.5 in trial 2. The second highest learning-rate interval contains all possible learning rates that would result in an associative strength equal to or larger than the threshold in trial 3, and smaller than the threshold in trial 2, and so forth. The lowest learning-rate interval contains all possible learning rates that would result in an associative strength equal to or larger than threshold in trial  $L+1$  and smaller than the threshold in trial  $L$ . (The lower interval limit of the lowest interval was set to  $-\ln(1-0.5)/L$  instead of zero in order to make the minimal resolution explicit.)

For estimating  $p(\epsilon_i)$  we use a simple search and count algorithm. For each behavioral sequence in a given data set we search for the most likely learning-rate interval. The probability distribution  $p(\epsilon_i)$  then equals the histogram of most likely counts for each interval, normalized by the total number of animals in a data set.

### The two-state hidden Markov model

The two-state hidden Markov model hypothesizes that at each trial  $t$  an animal can occupy one of two possible hidden states,  $k=1$  or  $k=2$  (see Bishop 2006 for a textbook account of hidden Markov models). Each of the two states is characterized by two probabilities for showing ( $x_t=1$ ) or not showing ( $x_t=0$ ) a CR, as well as by two probabilities for remaining in the same state or making a transition to the respective other state (see Fig. 4C). In total, the model comprises 10 parameters: two a priori state probabilities  $P(k_1)$ , four transition probabilities between states  $P(k_t|k_{t-1})$ , and four observation probabilities  $P(x_t|k_t)$  for observing the behavioral outcome  $x_t$  stemming from state  $k_t$ . Five of these 10 parameters are independent and were estimated by the Baum-Welch algorithm (Welch 2003) implemented in the hidden Markov model toolbox by Kevin Murphy for Matlab (University of British Columbia, Canada). Parameter estimation was insensitive to the choice of initial parameters for the data sets 1–15.

### Cross-validation

For each data set we computed the average log likelihood of the three models after 50 rounds of cross-validation. At each round, the data was split into four sets of equal size, each of which was then used once for testing and three times for training.

Data analysis and modeling was performed in Matlab (The MathWorks).

### Acknowledgments

We thank Jan Sölter, Bertram Gerber, Björn Brembs, and Paul Szyszka for helpful discussions and Chris Häusler for comments on the manuscript. This work was supported by the Bundesministerium für Bildung und Forschung (BMBF) through grant 01GQ0941 to D.E., R.M., and M.P.N. within the Bernstein Focus Neuronal Basis of Learning (BFNL). E.P. was funded by the Deutsche Forschungsgemeinschaft (DFG) in part through the Collaborative Research Center SFB 618 Theoretical Biology and a stipend

within the Research Training Group 1589 Sensory Computation in Neural Systems.

**Author contributions:** Research idea, data analysis, modeling, and manuscript by E.P. and M.P.N. Experimental data by N.K.C., L.M., V.A., J.F., K.B.G., N.S., D.E., and R.M.

### References

- Behrends A, Scheiner R. 2010. Learning at old age: A study on winter bees. *Front Neurosci* **4**: 15. doi: 10.3389/fnbeh.2010.000115.
- Bishop CM. 2006. *Pattern recognition and machine learning*. Springer, London, UK.
- Bitterman ME, Menzel R, Fietz A, Schäfer S. 1983. Classical conditioning of proboscis extension in honeybees (*Apis mellifera*). *J Comp Psychol* **97**: 107–119.
- Brown S, Heathcote AJ. 2003. QMLE: fast, robust, and efficient estimation of distribution functions based on quantiles. *Behav Res Methods Instrum Comput* **35**: 11–21.
- Chabaud M, Preat T, Kaiser L. 2010. Behavioral characterization of individual olfactory memory retrieval in *Drosophila melanogaster*. *Front Behav Neurosci* **4**: 192. doi: 10.3389/fnbeh.2010.00192.
- Cousineau D, Hélie S, Lefebvre C. 2003. Testing curvatures of learning functions on individual trial and block average data. *Behav Res Methods Instrum Comput* **35**: 493–503.
- Durstewitz D, Vitztoz NM, Floresco SB, Seamans JK. 2010. Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron* **66**: 438–448.
- Eisenhardt D. 2006. Learning and memory formation in the honeybee (*Apis mellifera*) and its dependency on the cAMP-protein kinase A pathway. *Anim Biol* **56**: 259–278.
- Estes WK. 2002. Traps in the route to models of memory and decision. *Psychon Bull Rev* **9**: 3–25.
- Felsenberg J, Gehring KB, Antemann V, Eisenhardt D. 2011. Behavioural pharmacology in classical conditioning of the proboscis extension response in honeybees (*Apis mellifera*). *J Vis Exp*. doi: 10.3791/2282.
- Gallistel CR, Fairhurst S, Balsam P. 2004. The learning curve: Implications of a quantitative analysis. *Proc Natl Acad Sci* **101**: 13124–13131.
- Gerber B, Wüstenberg D, Schütz A, Menzel R. 1998. Temporal determinants of olfactory long-term retention in honeybee classical conditioning: Nonmonotonic effects of the training trial interval. *Neurobiol Learn Mem* **69**: 71–78.
- Giurfa M. 2007. Behavioral and neural analysis of associative learning in the honeybee: A taste from the magic well. *J Comp Physiol [A]* **193**: 801–824.
- Hadar R, Menzel R. 2010. Memory formation in reversal learning of the honeybee. *Front Neurosci* **4**: 186. doi: 10.3389/fnbeh.2010.00186.
- Hanson SJ, Killeen P. 1981. Measurement and modeling of behavior under fixed-interval schedules of reinforcement. *J Exp Psychol Anim Behav Process* **7**: 129–139.
- Krechevsky I. 1932. "Hypotheses" in rats. *Psychol Rev* **39**: 516–532.
- Menzel R. 1999. Memory dynamics in the honeybee. *J Comp Physiol A* **185**: 323–340.
- Menzel R, Manz G, Greggers U. 2001. Massed and spaced learning in honeybees: The role of CS, US, the intertrial interval, and the test interval. *Learn Mem* **8**: 198–208.
- Menzel R, Brembs B, Giurfa M. 2007. Cognition in invertebrates. In *Evolution of nervous systems in invertebrates*, 403–422. Academic Press, Oxford, UK.
- Quinn WG, Harris WA, Benzer S. 1974. Conditioned behavior in *Drosophila melanogaster*. *Proc Natl Acad Sci* **71**: 708–712.
- Rath L, Giovanni Galizia C, Szyszka P. 2011. Multiple memory traces after associative learning in the honey bee antennal lobe. *Eur J Neurosci* **34**: 352–360.
- Rehder V. 1987. Quantification of the honeybees proboscis reflex by electromyographic recordings. *J Insect Physiol* **33**: 501–507.
- Rescorla R, Wagner A. 1972. A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: Current research and theory*, 64–99. Appleton-Century-Crofts, New York.
- Restle F. 1965. Significance of all-or-none learning. *Psychol Bull* **64**: 313–325.
- Roussel E, Sandoz J-C, Giurfa M. 2010. Searching for learning-dependent changes in the antennal lobe: Simultaneous recording of neural activity and aversive olfactory learning in honeybees. *Front Neurosci* **4**: 1–12.
- Sandoz JC, Roger B, Pham-Delègue MH. 1995. Olfactory learning and memory in the honeybee: Comparison of different classical conditioning procedures of the proboscis extension response. *C R Acad Sci III* **318**: 749–755.

- Scheiner R, Page RE, Erber J. 2004. Sucrose responsiveness and behavioral plasticity in honey bees (*Apis mellifera*). *Apidologie* **35**: 133–142.
- Schwärzel M, Müller U. 2006. Dynamic memory networks: Dissecting molecular mechanisms underlying associative memory in the temporal domain. *Cell Mol Life Sci* **63**: 989–998.
- Smith BH, Menzel R. 1989a. An analysis of variability in the feeding motor program of the honey bee: The role of learning in releasing a modal action pattern. *Ethology* **82**: 68–81.
- Smith BH, Menzel R. 1989b. The use of electromyogram recordings to quantify odourant discrimination in the honey bee, *Apis mellifera*. *J Insect Physiol* **35**: 369–375.
- Stollhoff N, Menzel R, Eisenhardt D. 2005. Spontaneous recovery from extinction depends on the reconsolidation of the acquisition memory in an appetitive learning paradigm in the honeybee (*Apis mellifera*). *J Neurosci* **25**: 4485–4492.
- Takeda K. 1961. Classical conditioned response in the honey bee. *J Insect Physiol* **6**: 168–179.
- Welch L. 2003. Hidden Markov models and the Baum-Welch algorithm. *IEEE Information Theory Society Newsletter* 53.

Received April 8, 2011; accepted in revised form August 30, 2011.