Reinforcement Expectation in the Honey Bee (*Apis mellifera*): Downshifts in Reinforcement Show Conditioned Inhibition

Shawn Mahoney[1*], Jay Hosler[2*], and Brian H Smith[1]

1 School of Life Sciences, Arizona State University, Tempe, AZ 85287-4501

2 Department of Biology, Juniata College, Huntingdon, PA 16652

*Co-equal first authors

Correspondence concerning this article should be addressed to brian.h.smith@asu.edu.

This work was supported by awards the BHS from the National Institutes of Health (GM 113967), the National Science Foundation (2014217) NeuroNex program, the National Science Foundation (2113179) CRCNS US-France (ANR) collaborative research program, and the Department of Energy (SC0021922) US-German (BMBF) collaborative research program.

ABSTRACT

When animals learn the association of a Conditioned Stimulus with an Unconditioned Stimulus, later presentation of the CS invokes a representation of the US. When the expected US fails to occur, theoretical accounts predict that conditioned inhibition can accrue to any other stimuli that are associated with this change in the US. Empirical work with mammals has confirmed the existence of conditioned inhibition. But the way it is manifested, the conditions that produce it, and determining whether it is the opposite of excitatory conditioning, are important considerations. Invertebrates can make valuable contributions to this literature because of the well-established conditioning protocols and access to the central nervous system for studying neural underpinnings of behavior. Nevertheless, conditioned inhibition has yet to be conclusively demonstrated in invertebrates. Here we evaluate the role of the unconditioned stimulus (US) in producing conditioned inhibition by using Proboscis Extension Response conditioning of the honey bee (*Apis mellifera*). Specifically, using variations of a 'feature-negative' experimental design, we employ downshifts in US intensity relative to US intensity used during initial excitatory conditioning, to show that an odorant in an odor-odor mixture can become a conditioned inhibitor. We argue that alternative interpretations to conditioned inhibition are unlikely. However, we show variation across individuals in how strongly they show Conditioned Inhibition, with some individuals possibly revealing a different means of learning about changes in reinforcement. We discuss how more work can reveal more details of conditioned inhibition, how it can be extended to investigate how it is encoded in the CNS, and how the work can generalize to other insect models. On particular, work like this will be important as more is revealed of the complexity of the insect brain from connectome projects.

INTRODUCTION

Animal studies of Pavlovian and operant conditioning have revealed important, generalizable principles of reinforcement learning (Sutton and Barto, 2018). Many studies have now shown that animals assess the informative value of stimuli and then update behavioral decisions based on that assessment every time these stimuli occur. At its most fundamental level, associative (Pavlovian) conditioning involves pairing of a 'neutral' Conditioned Stimulus (CS) with an important, salient Unconditioned Stimulus (US), such as food as an appetitive example or painful stimuli as aversive examples (Rescorla, 1988). The predictive nature of a CS produces a conditioned behavior that prepares animals for the encounter with the US. These principles have been expressed in simple mathematical equations aimed at explaining how conditioning establishes associative strength V_A to a CS_A, which then is the basis for anticipatory behavior (Bush and Mosteller, 1951; Rescorla and Wagner, 1972). Although there are important differences between the different models, such as whether and how attention contributes to conditioning (Lubow, 1989a; Mackintosh, 1975), these mathematical formulations share two important principles (Pearce and Hall, 1980; Rescorla, 1988). First, the equations assume that the change in associative strength ΔV can be either positive when the CS predicts the US will occur**,** or under at least some conditions negative when the CS predicts that an expected US will *not-occur*. Second, they all propose that ΔV on a given trial is proportional to the maximum amount of associative strength that a US can support (λ) minus an expectation of the summed total associative strength to all stimuli present – including contextual stimuli - that has already accrued regarding that US ($V_\Sigma=V_A+V_B+\ldots V_N$). Thus, associative strength V increases with each trial as long as $\lambda>V_\Sigma$, i.e. as long as the US that is encountered is stronger than predicted by V_Σ (a positive prediction error). Once the US is fully predicted ($\lambda=V_\Sigma$) associative strength stops being incremented, resulting in an asymptotic acquisition curve.

Behavioral studies have identified important manifestations of *reinforcement expectation* based on this type of model (Dunsmoor et al., 2015; Pearce and Hall, 1980; Rescorla and Wagner, 1972; Sutton and Barto, 2018). When an animal experiences CS_A after it has been associated with reinforcement, CS_A evokes a prediction of the US that is proportional to V_Σ. This amounts to an expectation that a US at least of value V_Σ will occur. When a second CS_X is combined with CS_A, and the same reinforcement is presented, animals typically learn less about CS_X than they normally would, which is called *blocking* (Mackintosh, 1983). Additionally, it is important to establish what happens on trials when the US fails to occur despite being predicted by CS_A, that is, when there is a surprising reward *omission* (Papini, 2003). In particular, when a second CS_X is present in compound (i.e. CS_{AX}) during reward omission, *conditioned inhibition to CS_X* may ensue

(Rescorla, 1969). Essentially, on such CS_{AX} trials $\lambda=0$ because the US is absent. However, $V_{\Sigma}>0$ because of the presence of CS_A. Therefore $\lambda-V_{\Sigma}$ is negative, reflecting that the US the animals encounter is less strong than what they expected. This *negative prediction error* serves as a teaching signal to establish a negative prediction for CS_X, i.e., the prediction that the US will not-occur. This is why CS_X becomes a *conditioned inhibitor*.

Revealing whether and under what conditions an animal develops and expresses Conditioned Inhibition can be critical for understanding how that animal is capable of representing information about reinforcement, and in particular which of the several models of conditioning apply (e.g. whether there is a representation akin to V_{Σ} or whether CS_{AX} is perceived as a different stimulus altogether). Moreover, showing the existence of Conditioned Inhibition will be important for understanding how that capability is supported by neural connectivity in the brain.

However, it is not trivial to unambiguously demonstrate Conditioned Inhibition (Papini and Bitterman, 1993; Sosa and Ramírez, 2019). Minimally, a conditioned inhibitor should pass two types of tests (Rescorla, 1969). First, once a stimulus (CS_X) has taken on inhibitory properties, subsequent excitatory conditioning should proceed more slowly than it does with other stimuli that do not carry inhibitory properties. A lack of retardation, or even its opposite, better excitatory conditioning of CS_X, might indicate that more attention is devoted to the processing of CS_X at the expense of other stimuli that are present. Second, CS_X should pass a summation test. When CS_X is presented in compound with a second CS that itself is capable of releasing a response, that response to the compound should be diminished by the inhibitory properties of CS_X.

The first step in this process is to demonstrate behavior that is consistent with Conditioned Inhibition. Although the existence of reward expectations has been well documented in vertebrates (Sosa and Ramírez, 2019), it has not been conclusively demonstrated via retardation and summation tests in invertebrates in spite of their widespread use as models for reinforcement learning. Using the honey bee (*Apis mellifera*), Chandra (Chandra et al., 2010) showed that unreinforced presentation of an odor X – as it would occur in an AX compound without reinforcement but where A predicts reinforcement, reliably produces retardation of acquisition. In other experiments, a reduction in expected reinforcement reduced learning of a target CS below that expected given the level of reinforcement (Smith, 1997). Furthermore, backward pairing (Hellstern et al., 1998) and peak shifts (Fernandez et al., 2009) produce inhibition-like effects. These experiments seem to rule out attention as a factor. However, the focus of these works was not on Conditioned Inhibition, so the results of summation tests were not evaluated.

Here we use variants of a feature-negative (A+|AX-) experimental design to test more directly for evidence of negative summation consistent with conditioned inhibition in the honey bee PER procedure. We show that when associated with omission of expected reinforcement, X can reduce the response to another excitatory stimulus. We also show and discuss that individuals differ in expression of Conditioned Inhibition, which parallels other studies of individual differences in learning performance in honey bees, and it brings into focus how potentially multiple conditioning phenomena can affect behavior. Finally, use of more complex and well controlled conditioning protocols to study learning should help map these phenomena onto the unexpected complexity of the fruit fly brain connectome (Li et al., 2020), and in particular it may reveal whether and how structures like the mushroom bodies are involved in reward expectation.

METHODS

Subjects

The subjects were honey bee workers obtained from colonies maintained out-of-doors in a genetically closed population (Page et al., 1985). For the first experiment, subjects were collected from outdoor colonies in the morning between 9 and 10 am and set-up in restraining harnesses according to published procedures (Smith and Burden, 2014). After a 15 minute period to acclimate to the harness, subjects were fed 0.4μl of a 1.25M sucrose solution, and they then remained undisturbed for 2 hours before the start of the conditioning procedure. All three experiments used a within-subjects design for testing.

Apparatus

Conditioning and testing were conducted in a 13 cm x 14.5 cm x 16.5 cm conditioning station illuminated with a fiber optic light source. The station had an open front, top and floor. The remaining sides were constructed of clear Plexiglas. To evacuate odors during conditioning, a 9.5 cm diameter hole in the back of the chamber was connected to an exhaust hood via a dryer tube. The floor of the chamber consisted of a single strip of Plexiglas connecting the left and right sides of the conditioning station. In the center of this strip was a small Plexiglas peg on which harnessed subjects were placed during conditioning.

Subjects were positioned to face an odor cartridge at the front of the chamber. Odor cartridges were made fresh daily by placing a strip of filter paper containing 3μl of pure odorant into a glass syringe. The open tip of the odor cartridge was positioned in front of and facing the subject. The ground glass fitting was connected to a three-way valve system that regulated the flow of air being

generated by an aquarium air pump. Administration of the CS (odor) was controlled by computer or Arduino-based logic controller. When triggered, the computer operated a solenoid valve, thus shunting air-flow away from an exhaust tube and through the odor cartridge. The odor-saturated air was ejected into the air stream that moved across the subject's antenna. Just behind the conditioning stage and out of the subject's view was a small LED used to signal the onset of odor delivery on videotape (see below). The US (0.4 µl of either 1.25 M or 0.50 M sucrose-water solution was administered using a Gilmont microliter syringe.

At the start of each trial the subject was placed into the conditioning station 20 seconds prior to the onset of the odor CS. It was also left in the conditioning station 30 seconds after odor exposure. The inter-trial interval for each subject was fixed at 6 min. A video camera was positioned above the conditioning stage to record each subject's response during the test phase of the experiment.

Procedure

Proboscis extension conditioning

In this protocol (Smith and Burden, 2014), an odorant (CS) is blown across the subject's antenna for 4 seconds. Three seconds after onset of the odor CS, a 0.4µl drop of sucrose-water (US) is lightly touched to the antenna. This stimulates sucrose taste receptors in the antenna and triggers the proboscis extension response. The subject is then allowed to consume the entire droplet, which always occurs within 2 sec. After 3-4 trials most subjects extend their proboscis to odor prior to sucrose reinforcement.

Experiment 1

The first experiment was divided into two acquisition phases during which six subjects were conditioned each day, so the intertrial interval was fixed at 6 min. In the initial (Pretraining) phase of all experiments, subjects were equivalently conditioned to six trials each of odor A and odor B presented in a pseudorandomized sequence (ABBABAABABBA). Each odor was appetitively reinforced, but with different sucrose concentrations. Odor A was reinforced with high (++; 1.25 M) sucrose and odor B was reinforced with low (+; 0.50 M) sucrose.

After pretraining, the first experiment consisted of a second acquisition phase followed by a test phase. For compound conditioning in the second phase all subjects received 12 trials. Subjects were exposed to odor A followed by High sucrose (A++ condition) on six of the trials. This condition was a continuation of the A++ condition from the preconditioning phase. On the remaining six trials, subjects were exposed to the compound of A and X followed by Low sucrose (AX+). Thus, X was reinforced with low sucrose in a context that predicted High sucrose

reinforcement. The two trial types were presented in the pseudo-randomized order used for pretraining.

After the second phase each subject was tested for their response to the two odor compounds BX and BN without the presence of reinforcement. Responses to odor B in combination with odor X were compared to responses to odor B when it was mixed with the novel odor N. The odors used for A, B, X, and N, and the order of presentation of BX and BN during testing, were counterbalanced over days of the experiment.

Experiments 2 and 3

Each subject was exposed to 32 acquisition trials, which was comprised of 8 trials each with A++, B+, AX- and CY- in a pseudorandomized sequence across trials. The major qualitative differences between this experiment and those reported above were: (1) that trials with A++ and AX- were pseudorandomly interspersed instead of presented in successive phases; and (2) the lack of reinforcement of the AX compound instead of reinforcement with a diminished US.

Following the acquisition trials, each subject was tested on unreinforced trials with BX and with BY in a random sequence across subjects. Odors 1-hexanol, 1-octanol, 2-hexanone and 2-octanone were counterbalanced as A, C, X and Y. There are 24 different combinations of 4 odorants; therefore, each of the 24 subjects was conditioned to a different pattern of A, C, X and Y. The fifth odorant (B), which was a common test stimulus, was the same for all subjects (geraniol).

Data analysis

After conditioning was complete all subjects received a nonreinforced presentation of odors BX and BN presented in a randomized order across days. All responses were video recorded for offline analysis. We then calculated the duration of proboscis extension during the 4 second presentation of odor X and for 20 seconds after the termination of odor presentation (Smith, 1997; Smith, 1998). Statistical analysis of the response duration was used for all hypothesis testing (Smith, 1998; Smith and Menzel, 1989). Statistical analyses were accomplished with paired comparison t-tests (Sokal and Rohlf, 1994). Means and standard errors are shown in the figures along with symbols indicating individual durations for each type of stimulus.

We present acquisition data for descriptive reasons. But we do not provide statistical tests of acquisition data, primarily because we do not use them for hypothesis testing. In some cases in particular, testing differences during acquisition trials, during which groups receive different treatment conditions, would confound treatment and testing conditions (Rescorla, 1988). Furthermore, it would be impossible to calculate duration of the conditioned response when its

expression would be influenced by the presentation of the US prior to offset of the CS. Given the similarity between the CR and the UR (Smith & Menzel, 1989), it would be impossible to distinguish them. We did not insert extinction trials to measure duration during acquisition, because it was not our intent to study differences in response to stimuli presented during the conditioning phases, and doing that might have attenuated the ability of a CS to induce inhibition.

Therefore, the only response measure we could use for acquisition was the percentage of subjects that responded to odor prior to presentation of the CS. That measure has been used in many studies of PER conditioning in the honey bee (Menzel, 1990), and when it reveals differences across treatment groups those differences are robust with almost any measure of the CR (Smith, 1998). But when it fails to reveal differences across treatments, use of a more parametric measure such as duration or latency can reveal differences in the response topology (Smith and Menzel, 1989).

RESULTS

Experiment 1: First test for negative summation

Subjects were initially conditioned in two phases. First, they to A++ and B+ on separate trials that were pseudorandomly interspersed with one another (Fig. 1A solid lines and filled symbols). For the second phase, subjects were conditioned to AX+ on trials that were similarly interspersed with continued presentation of A++ (Fig. 1A dashed lines and open symbols). The AX+ condition was designed to build inhibitory tendencies to X because of the decrease in reinforcement from A++ to AX+.

There was little difference between the preconditioning acquisition curves for odor A and odor B (Fig. 1A). In the compound conditioning phase, response to odor A alone (reinforced with High sucrose) remained high while the percent proboscis extension in the AX+ group (AX reinforced with Low sucrose) trended toward an overall decline. Since the strength of the US associated with A decreased between the preconditioning and compound conditioning phases in the AX+ trials, the decline in proboscis extension may already suggest an accrual of inhibition to odor X. If inhibition has accrued to X in the compound phase as a result of its association with a lower-than-expected US intensity, then this inhibitory effect should influence the subjects' relative response to the otherwise excitatory odor B when it is paired with X during testing.

We then tested subjects with a blend of B and X. Under normal circumstances, the addition of an odorant to a previously conditioned B would decrease the response to B because of overshadowing or interactions in sensory transduction processes (Smith, 1998). Therefore, we compare the response duration to B when it is combined with X to the response duration to B when it is combined with a novel odor (N). The inhibitory properties of X should decrease the response to B more than the addition of a novel odor N that is neutral.

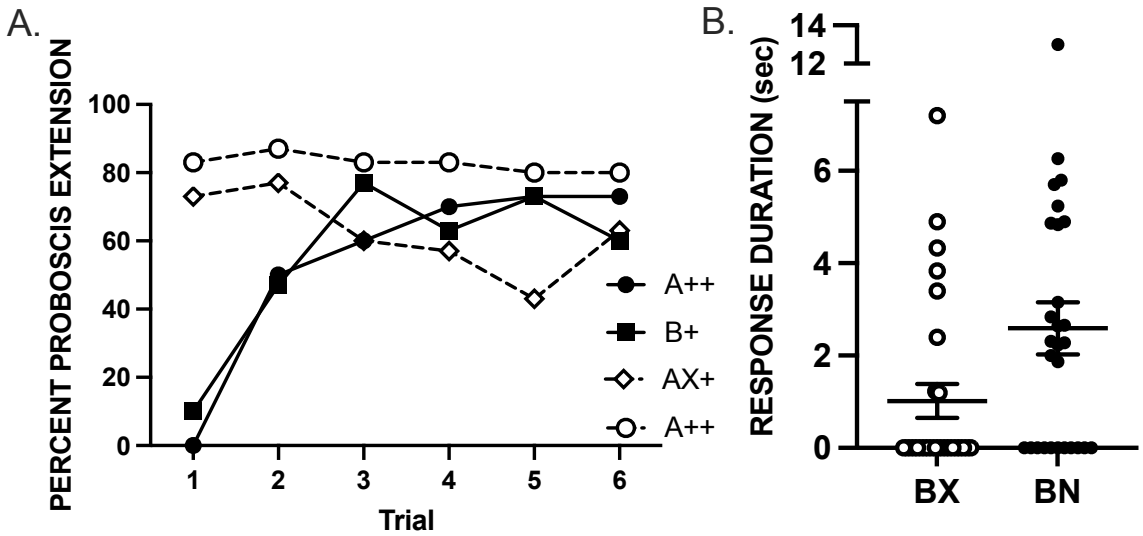

Figure 1. Summation test for conditioned inhibition. A, B, N and X were the odor Conditioned stimuli (CS). (A) Subjects (n = 30) were preconditioned over 12 conditioning trials (solid liens filled symbols). Six were performed with A++ and 6 with B+ presented in a pseudorandomized sequence. In the compound conditioning phase (dashed lines and open symbols), subjects were trained over six training trials each to A++ and AX+. (B) In the within-subjects design, all subjects were tested with the odor mixtures BN and BX. Symbols represent the mean (±SE) for each of the groups. Open and closed symbols represent response durations of individual subjects.

In the final summation test, the response to odor B should be suppressed in the BX mixture (because of its pairing with putative inhibitor X) relative to its response to odor B when it is paired to the novel odor N. The results indicate that this is indeed the case (Fig. 1B). The duration of proboscis extension response to BX was significantly shorter ($t(29) = 2.6$, $p < 0.01$ one-tailed paired comparison test) than that to BN. The distribution shown in Fig. 1B reflects a number of subjects that failed to respond to one or both odors, which are the points lying on the x-axis.

Experiment 2: A feature-negative test of Conditioned Inhibition

Most studies of conditioned inhibition have used a variation of a feature-negative design (Papini and Bitterman, 1993), which is a somewhat different procedure from that in our first

experiment. In a feature negative experiment we employ here, one odorant (A+) is paired with reinforcement, whereas a mixture of that odorant with a target odorant (AX-) is not reinforced at

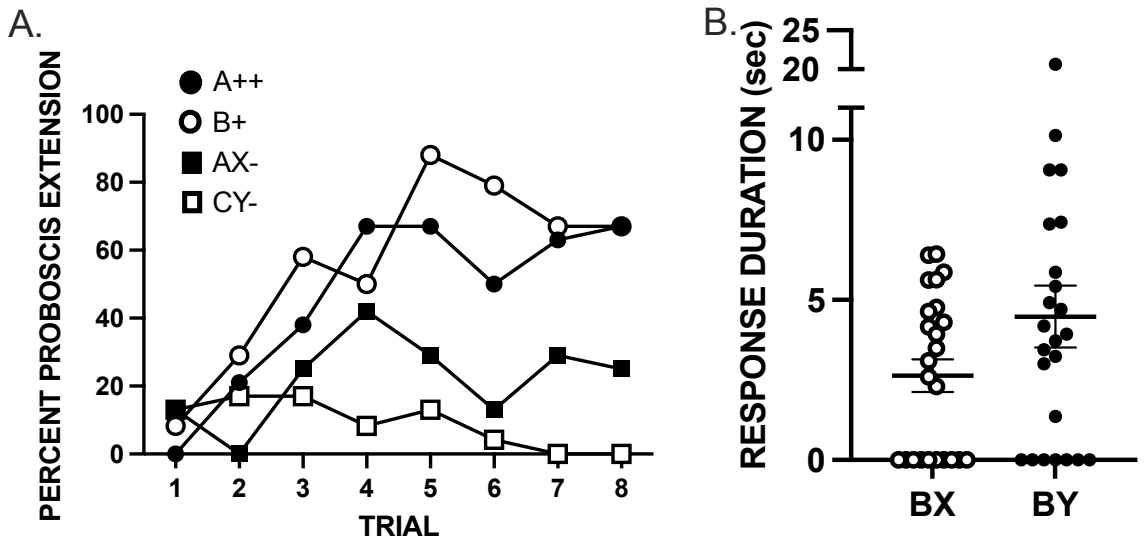

Figure 2. Test for conditioned inhibition using a feature-negative design. All subjects (n = 24) were exposed to 8 conditioning trials in each of the following 4 treatments: A++/B+/AX-/CY-. The 4 odorants used in this experiment were counterbalanced as A, X, C and Y. B was an excitatory test stimulus, and was always geraniol. (A) Acquisition to all 4 treatments. (B) Unreinforced summation tests with BX and BY. Horizontal and vertical, capped lines represent the mean (±SE) for each of the groups. Symbols are the durations for individual subjects.

all. We have used a similar procedure to demonstrate feature-negative conditioning in the honey bee (Chandra and Smith, 1998). This procedure produces a much more dramatic discrepancy between the reinforcement signaled by A and the actual outcome on trials when X is present. As in the first experiment, we used a single treatment group of 24 subjects. Each subject experienced four different types of trials that were presented intermixed in a pseudorandom order. Two types of trials involved pairing an odorant with either ++ or + reinforcement (A++ and B+) on separate trials. In this experiment we used B as an excitatory target odor against which to test for a summation effect. There were also two types of completely unreinforced trials (AX- and CY-). We predicted that inhibition would accrue to X on AX- trials because of the discrepancy between the signaled and actual outcomes. Furthermore, because unreinforced presentation of a CS may reduce its potential excitatory properties via latent inhibition (Chandra et al., 2010), we present a second unreinforced odor blend (CY-). If the inhibition that accrues to X is simply due to attentional decrement from nonreinforcement (latent inhibition), then there should be no difference between the response levels BX and BY in a subsequent summation test. On the other hand, if the lack of reinforcement in an otherwise excitatory context produces conditioned inhibition, then we would

expect that the response to BX would be less than that to BY. Furthermore, if afferent interaction in the AX compound reduces the response to BX, then there should be an equivalent reduction in BY because of afferent interaction in the CY condition.

Subjects showed rapid acquisition to A++ and to B+ (Fig. 2A), reaching an asymptotic level of response (circa 70%) by trials 4 or 5. The response level to AX- was intermediate, reaching a final response level of circa 25%. The response to CY-, the components of which were never reinforced, was lowest. The somewhat elevated level of response to AX- as compared to CY- could be due to incomplete excitatory generalization from A to the AX compound. Furthermore, inhibition to X could have also contributed to the response to AX- in the sense that it would decrease the response to AX relative to A.

Indeed, the summation test reveals that there is most likely a contribution from conditioned inhibition. The response to BX was significantly less than that to BY (Fig. 2B; $t(23) = 1.6$, $p < 0.05$ one-tailed test). Thus, odor X was capable of more significant negative summation when added to the conditioned excitor B than when Y was added.

Experiment 3: Individual differences in expression of Conditioned Inhibition

Figures 1 and 2 reflect mean test performance across a sample of 24 or 30 subjects, which is standard procedure for presenting learning performance in PER studies (Burden et al., 2016). However, mean values hide diversity of performance on any learning task, with some subjects performing well and others not showing the expected outcome. These differences are reflected in the distributions of individual duration values shown in Figs. 1B and 2B. Because of these individual differences, and its importance for a social colony (see Discussion; (Cook et al., 2019)), we decided to repeat the second experiment (Fig 2) with a focus on individual differences.

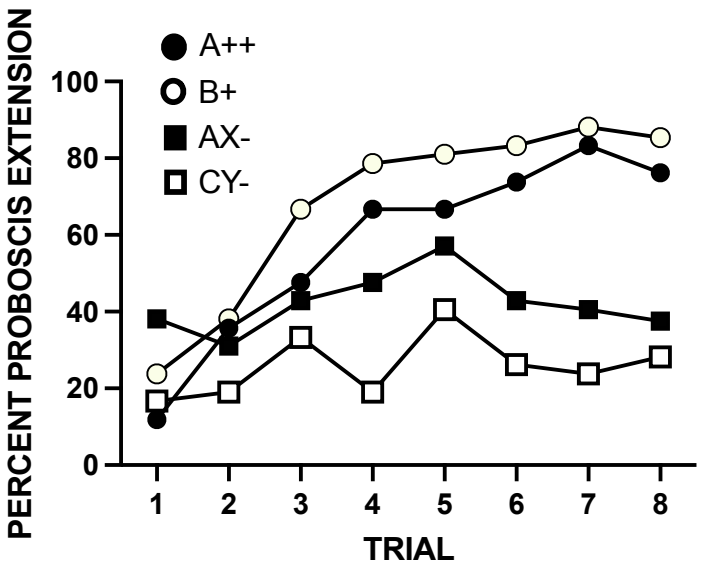

Figure 3. Acquisition phase for the replicate of experiment 2 using a feature-negative design. All subjects (n = 41) were exposed to 8 conditioning trials in each of the following 4 treatments: A++/B+/AX-/CY-. The 4 odorants used in this experiment were counterbalanced as A, X, C and Y. B was an excitatory test stimulus, and was always geraniol.

Acquisition across the four odors in this replicate was similar to acquisition in the first experiment (Figs. 2A). The response levels across trials to reinforced odors A++ and B+ were highest, reaching 70 to 80% of subjects responding or higher by the 5th or 6th trials. As before the response level to AX- was intermediate and the response to CY- was lowest.

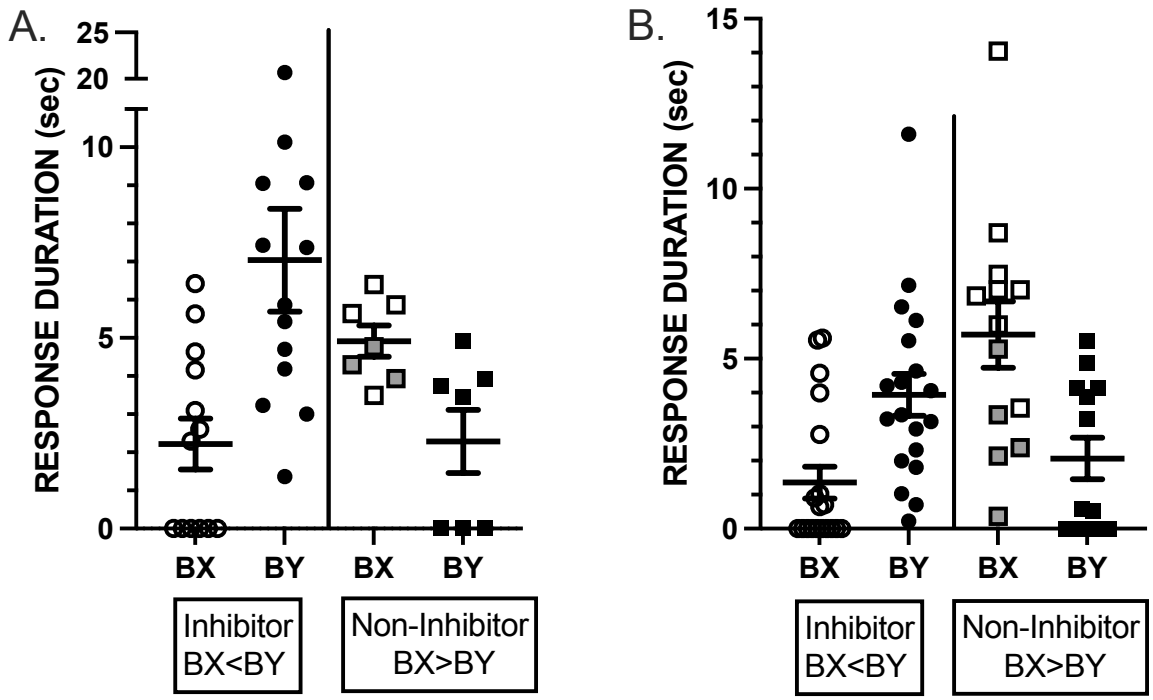

Figure 4. Summation tests for experiments 2 (A) and 3 (B). Plots show data from subjects that responded to one or both odors; subjects that failed to respond to either odor (4 in A and 9 in B) are not shown. Subjects in each experiment were classified as 'Inhibitor' if they showed a stronger response to BY than to BX. Otherwise, they were classified as Non-Inhibitor. Unreinforced summation tests with BX and BY. Horizontal and vertical, capped lines represent the mean (±SE) for each of the groups. Symbols are the durations for individual subjects. Shaded symbols for Non-Inhibitors are for subjects that responded to BX but not to BY.

In both experiments a majority of subjects that responded to one or both odors (13 of 20 and 19 of 32, resp) responded more strongly – i.e. with longer response durations - to BY than to BX (Fig. 4A,B; X^2 $p<0.01$), which is the pattern expected for Conditioned Inhibition. Data from these subjects are plotted as 'Inhibitors'. However, it is notable that a number of subjects in each experiment failed to show the pattern predicted by the Conditioned Inhibition hypothesis. First, some individuals failed to respond to either odor compound (4 and 9 in experiments 2 and 3, resp). Second, some subjects (7 of 20 and 13 of 32 in experiments 2 and 3, resp) showed shorter response durations to BY than to BX (right pairs of columns labelled 'Non-inhibitor'), which is

opposite to the prediction for Conditioned Inhibition. Obviously, because of the way they were selected, the differences within each group are statistically significant ($p<0.01$). The implications of these differences are further discussed below.

DISCUSSION

We have demonstrated that a component of a binary odor blend can develop properties of a conditioned inhibitor (Papini, 2003; Rescorla, 1969; Sosa and Ramírez, 2019). These properties develop when that component is associated with a lower level of reinforcement than that signaled by other components of the blend. However, there are several alternative interpretations that can account for a decrease in response to a target stimulus (Papini and Bitterman, 1993; Sosa and Ramírez, 2019). Those alternatives include modification of attention, afferent interaction, external inhibition, changes in generalized excitation, US habituation and contextual conditioning/blocking. It is important now to address these issues in order to determine the reliability of the claim for inhibition in our experiments.

Rescorla (Rescorla, 1969) has recommended a two-test strategy for revealing conditioned inhibition, which would serve as a necessary set of conditions for showing inhibition. A conditioned inhibitor would reveal retardation of acquisition when it is properly paired with reinforcement, and its inhibition would negatively summate with the excitatory properties of another stimulus that was associated with the same type of reinforcement used to produce the inhibition. If A+/AX- pairing changed attention to X, then X might pass one test but not the other. For example, if attention to X were enhanced because of the *surprising* omission of reinforcement in the A+/AX- condition, then acquisition would be expected to be more rapid to X in an A+/AX- group than in a control group (e.g., A+/X-). In this case X would fail the retardation test but might still decrease the response to an excitatory stimulus because of X's ability to monopolize attention. On the other hand, if the effect of nonreinforcement of X in the A+/AX- condition were to decrease attention to X via latent inhibition (Lubow, 1973), for example, then acquisition would be retarded, and X would pass the retardation test, but X might fail to reduce the excitatory properties of another stimulus leading to failure of the summation test. Only a conditioned inhibitor would be capable of passing both tests.

Interpretation of reduced response to BX in our experiments, relative to BN, because of increased attention to X is unlikely. If attention were a factor, then any modification of the US between the initial and compound trials will generate some form of surprise and hence enhanced attention to X (Mackintosh, 1983; Pearce and Hall, 1980). That is, either an up- or a down-shift in

reinforcement should be capable of generating a significant amount of learning about X on compound (AX) trials. However, when honey bees experience A++/AX++ training, they learn less about X relative to a control procedure in which A is not present during the compound phase (Smith, 1997). Furthermore, when they experience a downshift in reinforcement (A++/AX+), they respond less to X relative to the A++/AX+ group that did not experience the downshift. The attention hypothesis would predict the opposite. Moreover, this procedure showed one-trial blocking. Those results are inconsistent with the interpretation that attention to X is enhanced by a downshift in reinforcement used in our current experiments (Mackintosh, 1983; Rescorla and Holland, 1982).

Latent inhibition (Lubow, 1989b; Lubow, 1973) via retardation of acquisition also does not account for our results. Latent inhibition in the honey bee PER paradigm requires many more (~20) unreinforced exposures to a CS (Chandra et al., 2010) than subjects were exposed to in our experiment. Nevertheless, we attempted to control for this effect in experiments 2 and 3 by exposing subjects to the CY- condition. If latent inhibition decreased response to B in the AB- condition, then it should have done so equally to Y in the CY- condition. But the test response to BX was lower than to BY, which argues at least for an additional source of inhibition to X.

Afferent interaction might reasonably be expected to be stronger in intramodal compounds than in intermodal compounds (Kehoe et al., 1994). We attempted to control for afferent interaction by the addition of a novel odorant during test phase, as we previously have done for studies of blocking (Smith, 1996; Smith, 1997; Smith and Cobey, 1994; Thorn and Smith, 1997). If we had simply tested BX versus B, then the decrease would be interpretable in terms of greater interaction in the compound, which occurs in mixtures of these same odorants used above (Smith, 1998). Furthermore, we either partially or completely counterbalanced odorants in all of our experiments. It is therefore unlikely that the summation effect of X would arise due to greater afferent interaction with B than in the control groups.

Several previous attempts have been made to specifically study conditioned inhibition in honey bees, but the results were mixed. Some early experiments using different types of manipulations (Bitterman et al., 1983), including backward pairing (Hellstern et al., 1998) and peak shift (Fernandez et al., 2009), provided an indication of inhibition, but the experiments left open alternative interpretations. Couvillon et al. (Couvillon et al., 1999) failed to find evidence for conditioned inhibition in freely flying honey bees. Yet, use of modified, and perhaps more sensitive, experimental designs detected evidence of conditioned inhibition (Couvillon et al., 2003). Using restrained subjects in the PER paradigm provided indirect evidence for inhibition (Couvillon et al., 2005). Yet our experiments using PER reported here support the interpretation,

over other alternative possibilities, that conditioned inhibition can be found in honey bees. The design we adopted for experiments 2 and 3 was adopted from Couvillon et al. (Couvillon et al., 1999). Therefore, differences across studies do not appear to be attributable to use of freely flying versus restrained bees, or to the use of intra- versus inter-modal stimulus mixtures.

A major difference between our study reported here and, other studies using PER conditioning of restrained subjects (Couvillon et al., 2005) involves aspects of experimental design and use of different response measures. Our designs employed longer phases of excitatory conditioning – 18 or 32 trials versus 9. A similar difference can account for differences in detection of latent inhibition (Bitterman et al., 1983; Chandra et al., 2010). Furthermore, studies of PER conditioning have mostly used the percentage of subjects in a treatment group that responded to the CS. This measure is frequently sensitive enough to detect differences between treatment groups. But failure to find differences, particularly when a large percentage of subjects respond, cannot be interpreted as unambiguous evidence of a lack of treatment effect. Although the response itself looks simple, detailed electromyographic (Smith and Menzel, 1989) and video analyses like we have done here reveal several different features of response topology (e.g. response latency and duration) that differentiate treatment groups (Smith, 1997), even when no differences are evident in response probability.

One other difference between our current study and others is that we show individual differences in manifestation of Conditioned Inhibition, which are present in all studies of PER conditioning. Several studies of PER in honey bees have shown consistent individual differences in performance in several different learning protocols (Finke et al., 2023; Scheiner et al., 2017), which also broadly correlates to sensitivity to sugar as well as to other foraging preferences (Scheiner et al., 2001). Furthermore, much of this variation in how worker honey bees perform on any learning protocol can be attributed to genotype (Chandra et al., 2000; Chandra et al., 2001; Latshaw et al., 2023). Some workers perform any learning task well, whereas others do not. Because a queen mates with up to 20 drones (males), any colony will contain a mixture of several different patrilines (Page, 2013). This genetic diversity within any colony is important for the fitness of the colony as it performs foraging tasks (Cook et al., 2020; Mosqueiro et al., 2017). However, this diversity must be taken into account in any study of conditioning, because a small sample (~20 or 30) workers from a large population (up to 100,000) may not reflect the genetic diversity present in the large population.

We confirm here that the groups of subjects selected for experiments 2 and 3 show differences across individuals in how the behavior is manifested. The majority of individuals responded more strongly to the BY compound than to the BX compound, which is the prediction for Conditioned

Inhibition. In contrast, a minority of individuals in each experiment responded more strongly to BX than to BY, which is the opposite prediction from Conditioned Inhibition. This pattern (BX>BY) could arise if some individuals associated X with excitatory conditioning of A++, as would be expected for Second Order conditioning (Mackintosh, 1983), for example. In fact, the procedure we have employed would also be employed to study second order conditioning. In that case, stronger response to the BX compound could occur via *positive* summation instead of the negative summation predicted for Conditioned Inhibition.

These individual differences in our studies leave open possibilities for further investigation. For example, do some individuals in a colony show excitatory second order conditioning over Conditioned Inhibition, or vice versa, when a stimulus is paired with a conditioned excitor? Does excitatory second order conditioning drive responding over the first few conditioning trials and then switch to Conditioned Inhibition after several trials? In this case, some individuals may make this switch earlier than others, which would account for individual differences in our study. Within colony genetic differences have been shown, for example, for Latent Inhibition and Reversal Learning (Chandra et al., 2010; Ferguson, 2001), and could carry over to other forms of conditioning. Our results emphasize that future studies need to describe individual differences in learning performance and test how other factors – age, experience, behavioral caste, genetic background – affect expression of the trait. At the very least, care must be taken to control for and standardize genetic background, which is not often done in studies with honey bees, but is standard practice in studies with fruit flies and rodents.

Taken together, studies now show enough evidence that conditioned inhibition can be studied in honey bees, and now also in molluscs (Acebes et al., 2012). However, much needs to be done to more fully evaluate how conditioned inhibition is specifically manifested in behavior. For example, does conditioned inhibition show extinction, like in excitatory conditioning? What is the role of contextual cues in generating and releasing memory for conditioned inhibition? Are there differences in manifestation of conditioned inhibition between inter- and intra-modal mixtures, as were used in our study? Answers to these questions will be essential for understanding how different forms of learning are integrated into the ongoing decision processes involved in acquiring resources or avoiding threats (Bazhenov et al., 2013). Additionally, this information will interface with investigations into the neural mechanisms that underlie Conditioned Inhibition, such as have been initiated for latent inhibition (Locatelli et al., 2013) and blocking (Chen et al., 2014) in honey bees. These studies in insects such as honey bees and fruit flies, particularly if conditioned inhibition can be found in the latter, will help integrate the behavior into information now coming out of the fruit fly brain connectome projects (Li et al., 2020).

REFERENCES

Acebes, F., Solar, P., Moris, J. and Loy, I. (2012). Associative learning phenomena in the snail (Helix aspersa): conditioned inhibition. *Learn Behav* **40**, 34-41.

Bazhenov, M., Huerta, R. and Smith, B. H. (2013). A Computational Framework for Understanding Decision Making through Integration of Basic Learning Rules. *Journal of Neuroscience* **33**, 5686-5697.

Bitterman, M. E., Menzel, R., Fietz, A. and Schafer, S. (1983). Classical conditioning of proboscis extension in honeybees (*Apis mellifera*). *Journal of Comparative Psychology* **97**, 107-19.

Burden, C. M., Elmore, C., Hladun, K. R., Trumble, J. T. and Smith, B. H. (2016). Acute exposure to selenium disrupts associative conditioning and long-term memory recall in honey bees (Apis mellifera). *Ecotoxicol Environ Saf* **127**, 71-9.

Bush, R. R. and Mosteller, F. (1951). A mathematical model for simple learning. *Psychol Rev* **58**, 313-23.

Chandra, S. and Smith, B. H. (1998). An analysis of synthetic processing of odor mixtures in the honeybee (Apis mellifera). *Journal of Experimental Biology* **201**, 3113-21.

Chandra, S. B., Hosler, J. S. and Smith, B. H. (2000). Heritable variation for latent inhibition and its correlation with reversal learning in honeybees (Apis mellifera). *Journal of Comparative Psychology* **114**, 86-97.

Chandra, S. B., Hunt, G. J., Cobey, S. and Smith, B. H. (2001). Quantitative trait loci associated with reversal learning and latent inhibition in honeybees (Apis mellifera). *Behavior Genetics* **31**, 275-85.

Chandra, S. B., Wright, G. A. and Smith, B. H. (2010). Latent inhibition in the honey bee, Apis mellifera: Is it a unitary phenomenon? *Anim Cogn* **13**, 805-15.

Chen, e.-Y., Marachlian, E., Assisi, C., Huerta, R., Smith, B., Locatelli, F. and Bazhenov, M. (2014). Learning modifies odor mixture processing to improve detection of relevant components. *Journal of Neuroscience*.

Cook, C. N., Lemanski, N. J., Mosqueiro, T., Ozturk, C., Gadau, J., Pinter-Wollman, N. and Smith, B. H. (2020). Individual learning phenotypes drive collective behavior. *Proc Natl Acad Sci U S A* **117**, 17949-17956.

Cook, C. N., Mosqueiro, T., Brent, C. S., Ozturk, C., Gadau, J., Pinter-Wollman, N. and Smith, B. H. (2019). Individual differences in learning and biogenic amine levels influence the behavioural division between foraging honeybee scouts and recruits. *J Anim Ecol* **88**, 236-246.

Couvillon, P. A., Ablan, C. D. and Bitterman, M. E. (1999). Exploratory studies of inhibitory conditioning in honeybees (Apis mellifera). *J Exp Psychol Anim Behav Process* **25**, 103-12.

Couvillon, P. A., Bumanglag, A. V. and Bitterman, M. E. (2003). Inhibitory conditioning in honeybees. *Q J Exp Psychol B* **56**, 359-70.

Couvillon, P. A., Hsiung, R., Cooke, A. M. and Bitterman, M. E. (2005). The role of context in the inhibitory conditioning of honeybees. *Q J Exp Psychol B* **58**, 59-67.

Dunsmoor, J. E., Niv, Y., Daw, N. and Phelps, E. A. (2015). Rethinking Extinction. *Neuron* **88**, 47-63.

Ferguson, H. J., Cobey, S., Smith, B.H. (2001). Sensitivity to a change in reward is heritable in the honey bee, Apis mellifera. *Animal Behaviour* **61**, 527-534.

Fernandez, P. C., Locatelli, F. F., Person-Rennell, N., Deleo, G. and Smith, B. H. (2009). Associative conditioning tunes transient dynamics of early olfactory processing. *J Neurosci* **29**, 10191-202.

Finke, V., Scheiner, R., Giurfa, M. and Avarguès-Weber, A. (2023). Individual consistency in the learning abilities of honey bees: cognitive specialization within sensory and reinforcement modalities. *Anim Cogn* **26**, 909-928.

Hellstern, F., Malaka, R. and Hammer, M. (1998). Backward inhibitory learning in honeybees: a behavioral analysis of reinforcement processing. *Learn Mem* **4**, 429-44.

Kehoe, E. J., Horne, A. J., Horne, P. S. and Macrae, M. (1994). Summation and configuration between and within sensory modalities in classical conditioning of the rabbit. *Animal Learning & Behavior* **22**, 19-26.

Latshaw, J. S., Mazade, R. E., Petersen, M., Mustard, J. A., Sinakevitch, I., Wissler, L., Guo, X., Cook, C., Lei, H., Gadau, J. et al. (2023). Tyramine and its *Amtyr1* receptor modulate attention in honey bees (Apis mellifera). *Elife* **12**.

Li, F., Lindsey, J. W., Marin, E. C., Otto, N., Dreher, M., Dempsey, G., Stark, I., Bates, A. S., Pleijzier, M. W., Schlegel, P. et al. (2020). The connectome of the adult Drosophila mushroom body provides insights into function. *Elife* **9**.

Locatelli, F. F., Fernandez, P. C., Villareal, F., Muezzinoglu, K., Huerta, R., Galizia, C. G. and Smith, B. H. (2013). Nonassociative plasticity alters competitive interactions among mixture components in early olfactory processing. *Eur J Neurosci* **37**, 63-79.

Lubow, R. (1989a). Latent Inhibition and Conditioned Attention Theory. Cambridge University Press.

Lubow, R. (1989b). Latent Inhibition and Conditioned Attention Theory. Cambridge University Press: Cambridge University Press.

Lubow, R. E. (1973). Latent inhibition. *Psychological Bulletin* **79**, 398-407.

Mackintosh, N. (1983). Conditioning and Associative Learning: Clarendon Press.

Mackintosh, N. J. (1975). Blocking of conditioned suppression: role of the first compound trial. *J Exp Psychol Anim Behav Process* **1**, 335-45.

Menzel, R. (1990). Learning, memory, and 'cognition' in honeybees. In *Neurobiology of comparative cognition*, pp. 237-292. Hillsdale, N.J.: Lawrence Erlbaum.

Mosqueiro, T., Cook, C., Huerta, R., Gadau, J., Smith, B. and Pinter-Wollman, N. (2017). Task allocation and site fidelity jointly influence foraging regulation in honeybee colonies. *R Soc Open Sci* **4**, 170344.

Page, R. E. (2013). The Spirit of the Hive: The Mechanisms of Social Evolution: Harvard University Press.

Page, R. E., Erickson, E. H. and Laidlaw, H. H. (1985). Closed population honey bee breeding 4. The distribution of sex alleles with top outcrossing. *Journal of Apicultural Research* **24**, 38-42.

Papini, M. R. (2003). Comparative psychology of surprising nonreward. *Brain Behav Evol* **62**, 83-95.

Papini, M. R. and Bitterman, M. E. (1993). The two-test strategy in the study of inhibitory conditioning. *J Exp Psychol Anim Behav Process* **19**, 342-52.

Pearce, J. M. and Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol Rev* **87**, 532-52.

Rescorla, R. A. (1969). Pavlovian conditioned inhibition, vol. 72, pp. 77-94. Psychological Bulletin.

Rescorla, R. A. (1988). Behavioral studies of Pavlovian conditioning. *Annu Rev Neurosci* **11**, 329-52.

Rescorla, R. A. and Holland, P. C. (1982). Behavioral studies of associative learning in animals. In *Annual Review of Psychology*, vol. 33, pp. 265-308.

Rescorla, R. A. and Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical Conditioning II*, eds. A. H. Black and W. F. Prokasy), pp. 64-99: Appleton-Century-Crofts.

Scheiner, R., Page, R. E., Jr. and Erber, J. (2001). Responsiveness to sucrose affects tactile and olfactory learning in preforaging honey bees of two genetic strains. *Behavioural Brain Research* **120**, 67-73.

Scheiner, R., Reim, T., Søvik, E., Entler, B. V., Barron, A. B. and Thamm, M. (2017). Learning, gustatory responsiveness and tyramine differences across nurse and forager honeybees. *J Exp Biol* **220**, 1443-1450.

Smith, B. H. (1996). The role of attention in learning about odorants. *Biol Bull* **191**, 76-83.

Smith, B. H. (1997). An analysis of blocking in odorant mixtures: an increase but not a decrease in intensity of reinforcement produces unblocking. *Behav Neurosci* **111**, 57-69.

Smith, B. H. (1998). Analysis of interaction in binary odorant mixtures. *Physiology & Behavior* **65**, 397-407.

Smith, B. H. and Burden, C. M. (2014). A proboscis extension response protocol for investigating behavioral plasticity in insects: application to basic, biomedical, and agricultural research. *J Vis Exp*, e51057.

Smith, B. H. and Cobey, S. (1994). The olfactory memory of the honeybee Apis mellifera. II. Blocking between odorants in binary mixtures. *Journal of Experimental Biology* **195**, 91-108.

Smith, B. H. and Menzel, R. (1989). 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.

Sokal, R. R. and Rohlf, F. J. (1994). Biometry: The Principles and Practice of Statistics in Biological Research: W.H. Freeman.

Sosa, R. and Ramírez, M. N. (2019). Conditioned inhibition: Historical critiques and controversies in the light of recent advances. *J Exp Psychol Anim Learn Cogn* **45**, 17-42.

Sutton, R. S. and Barto, A. G. (2018). Reinforcement Learning: An Introduction. MIT Press, Cambridge MA.

Thorn, R. S. and Smith, B. H. (1997). The olfactory memory of the honeybee Apis mellifera. III. Bilateral sensory input is necessary for induction and expression of olfactory blocking. *Journal of Experimental Biology* **200**, 2045-55.