

LEARNING & MEMORY

Context and occasion setting in *Drosophila* visual learning

Björn Brembs and Jan Wiener

Learn. Mem. 2006 13: 618-628

Access the most recent version at doi:[10.1101/lm.318606](https://doi.org/10.1101/lm.318606)

References

This article cites 62 articles, 22 of which can be accessed free at:
<http://www.learnmem.org/cgi/content/full/13/5/618#References>

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](#)

Notes

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



Research

Context and occasion setting in *Drosophila* visual learning

Björn Brembs^{1,4} and Jan Wiener^{2,3}

¹Institute of Biology, Neurobiology, Freie Universität Berlin, Königin-Luise-Strasse 28/30, 14195 Berlin, Germany; ²University of Würzburg, Department of Genetics and Neurobiology, Biozentrum am Hubland, 97074 Würzburg, Germany

In a permanently changing environment, it is by no means an easy task to distinguish potentially important events from negligible ones. Yet, to survive, every animal has to continuously face that challenge. How does the brain accomplish this feat? Building on previous work in *Drosophila melanogaster* visual learning, we have developed an experimental methodology in which combinations of visual stimuli (colors and patterns) can be arranged such that the same stimuli can either be directly predictive, indirectly predictive, or nonpredictive of punishment. Varying this relationship, we found that wild-type flies can establish different memory templates for the same contextual color cues. The colors can either leave no trace in the pattern memory template, leading to context-independent pattern memory (context generalization), or be learned as a higher-order cue indicating the nature of the pattern-heat contingency leading to context-dependent memory (occasion setting) or serve as a conditioned stimulus predicting the punishment directly (simple conditioning). In transgenic flies with compromised mushroom-body function, the sensitivity to these subtle variations is altered. Our methodology constitutes a new concept for designing learning experiments. Our findings suggest that the insect mushroom bodies stabilize visual memories against context changes and are not required for cognition-like higher-order learning.

Πάντα ρεῖ καὶ οὐδὲν μένει—Everything flows, nothing stands still (Heraclitus). Rapid changes in environmental contingencies require flexible capacities through which organisms can come to expect biologically significant events (unconditioned stimuli, US) and modify the behavior in anticipation of those events if behavior is to remain adaptive; i.e., increase the probability of obtaining beneficial and avoiding harmful consequences (Sutton and Barto 1998; Dickinson and Balleine 2002). In a dynamic environment, some of the stimuli can predict the occurrence of single USs (conditioned stimuli, CS), others may indicate the nature of the CS–US contingency (occasion setters, OS) and again others may be present without having any relationship to the US whatsoever (context). Thus, in order to be able to form an accurate expectation of future USs, animals have to extract from the universe of sensory signals the actual predictors by separating them from nonpredictive stimuli. In principle, this can be achieved if only those sensory inputs that bear a temporal relationship to the reinforcer are taken as predictors (Wickens 1987).

Tethered *Drosophila* can be trained to avoid heat punishment (US). Predictors of heat punishment can be the behavior of the fly, a variety of stimuli or almost any combination of both (Wolf and Heisenberg 1991, 1997; Wolf et al. 1998; Ernst and Heisenberg 1999; Liu et al. 1999; Brembs and Heisenberg 2000, 2001; Heisenberg et al. 2001; Tang et al. 2004; Katsov and Clandinin 2006). In these paradigms, the fly is attached to a measuring device that transduces the fly's turning behavior (yaw torque) into an analog signal (Fig. 1). The signal can be used to establish any kind of behavioral consequence (Heisenberg et al. 2001). We have used the unique environmental control this set-up affords to highlight the role of the temporal relationship of initially neutral stimuli (context, CS, OS) and US, and its consequences for

the acquisition of predictive memory in wild-type and transgenic flies.

Take, for instance, differential conditioning of visual patterns (Wolf and Heisenberg 1991; Ernst and Heisenberg 1999; Brembs and Heisenberg 2000; Tang et al. 2004; Katsov and Clandinin 2006). In this paradigm, animals learn to avoid one visual pattern (e.g., an upright T) and to prefer another (e.g., an inverted T). All other stimuli remain constant throughout the experiment. Slightly modifying the experiment by changing the background color between training and test (e.g., from blue-green to blue or from blue-green to green or vice versa) does not disrupt performance (Liu et al. 1999). The color remains constant during training and thus the T-patterns are the sole reliable predictors of reinforcement—the colors fulfill the definition of context. Wild-type flies can generalize the pattern memory across certain contexts (context-independent memory), while flies with impaired mushroom-body function cannot (Liu et al. 1999). The pattern memory of the mushroom-body-impaired flies is thus context-dependent. Interestingly, context dependence is often presented as a costly or advanced brain capacity or feature, while context independence is often described as a failure of the brain to incorporate the context into the memory template (e.g., Law et al. 2004). It is curious that flies with impaired mushroom-body function should exhibit such a feature, while wild-type flies fail to do so. Is context dependence a feature or a failure of the brain? One explanation for the low learning scores in the transgenic or mushroom-bodyless flies may be that they are not able to perform the separation between patterns and colors (Liu et al. 1999). In this view, the exhibited context dependence is a failure to separate patterns from colors, forming a compound memory template. A second explanation may be that the flies with compromised mushroom-body function detect much more quickly than wild-type flies that there is no punishment at all in the new situation. This enhanced detection then abolishes the pattern preference in the mushroom-bodyless flies. In this view, the context dependence of the pattern memory is a feature of the brain, not present in wild-type flies and brought about by inhibiting

³Present address: Collège de France, LPPA 11, Place Marcelin Berthelot, 75231 Paris.

⁴Corresponding author.

E-mail bjoern@brembs.net; fax 49-308-385-5455.

Article is online at <http://www.learnmem.org/cgi/doi/10.1101/lm.318606>.




Figure 1. Flight simulator set-up and experimental schematics. (Left) The fly is flying stationarily in a cylindrical arena homogeneously illuminated from behind. The fly's tendency to perform *left* or *right* turns (yaw torque) is measured continuously and fed into the computer. The computer controls arena position, IR-laser (heat) activation and color of illumination according to the conditioning rules. (Right) Experimental schematics used in this study. Patterns and colors depict the wall of the cylinder surrounding the fly. Colored boxes indicate the four 90° quadrants. Red boxes in the example color/pattern schematics depict heated quadrants. Note that even though quadrants may be drawn in different colors here, the illumination of the entire arena is always changed. See Materials and Methods for a detailed description.

mushroom-body output. To test these hypotheses, we have designed a set of experiments in which the same stimuli are arranged to produce both context-dependent and independent pattern memory in wild-type flies (Fig. 2). To start exploring the biological basis of context dependence, we tested transgenic flies in some of these experiments (Fig. 4, below).

Results

Outline

We observed a shift in associative strength of the contextual color cues with increasing predictive relationship of the colors to the punishing heat beam (Fig. 2). The increase in the predictive relationship was conducted in five steps. The first step was to reproduce the original context generalization experiment by Liu et al. (1999) with wild-type flies and also with the *mb247* driver strain used here (driving tetanus-toxin expression in the mushroom bodies). In this case, the predictive relationship was minimal, as the color change only occurred once, and this single signal predicted a test period (Fig. 2A, Context generalization). In a second step, the background color change still only predicted test from training, but this training-test transition now occurred repeatedly instead of only once (Fig. 2B,C, Feature discrimination). In a third step, background colors were still changed between periods as before (i.e., temporally), but this time the color change did not simply indicate the training-test transitions, but indicated a reversal of the CS-US contingency; one color indicated that the upright T was punished (and the inverted T unpunished), while the other color indicated the reverse relationship (Fig. 2D, Temporal conditional discrimination). In a fourth step, the colors still indicated the nature of the CS-US contingency as in the step before, but instead of only controlling the patterns while the colors were switched between periods, the flies now had operant control over both colors and patterns (Fig. 2E, Operant spatial conditional discrimination). A yoked control established the importance of this additional operant component (Fig. 2F, Yoked conditional discrimination). At the same step, we also established classical conditional discrimination (Fig. 2G, Classical spatial conditional discrimination). In a fifth and final step, we set up the background colors as direct predictors (CS) of the US (Fig. 2H, Operant color discrimination learning). Thus, in

a battery of tests, we varied the predictive value of the same contextual color cues in a stepwise fashion from minimally predictive (Step 1) to indirectly predictive (Steps 2–4) to directly predictive (Step 5). Would the flies follow this scheme and shift their processing of the colors from generalization to discrimination?

To investigate the biological basis of context dependence and to provide a proof of concept for the neurobiological value of such a closed methodology, we tested flies with experimentally blocked mushroom-body output in both context generalization and two cases of conditional discrimination (Fig. 4, below).

Colors as context

At the first step, wild-type flies master the context generalization task (Fig. 2A, Context generalization), while the transgenic flies with blocked mushroom-body output fail to generalize the pattern memory across the contexts (Fig. 4A, below, Context generalization). This result corroborates and extends the results by Liu et al. (1999), which did not include the driver strain *mb247*. Thus, the wild-type animals did not reveal whether they had detected the context change and showed the conditioned pattern preference even in the new context. This result simultaneously corroborates previous findings (Brembs and Heisenberg 2001; Tang and Guo 2001) that flies can process patterns independently from colors and do not treat the two sets of stimuli as a compound. It needs to be pointed out that flies with impaired mushroom-body function are otherwise fairly normal and, for example, readily learn to discriminate the visual patterns operantly and classically (Wolf et al. 1998; Liu et al. 1999).

Colors in a feature discrimination task

In the context generalization experiment, the background illumination is only changed once, giving this change in the total stimulus situation minimal predictive value. In our second step, we increased the predictive value of this change by increasing the number of color changes together with the application of the reinforcer. Switching between training and test periods every minute reduces the final test score only if training and test phases are characterized by different background colors (see Fig. 2B, Feature discrimination and Fig. 2C, Constant context). The flies have learned to conditionally

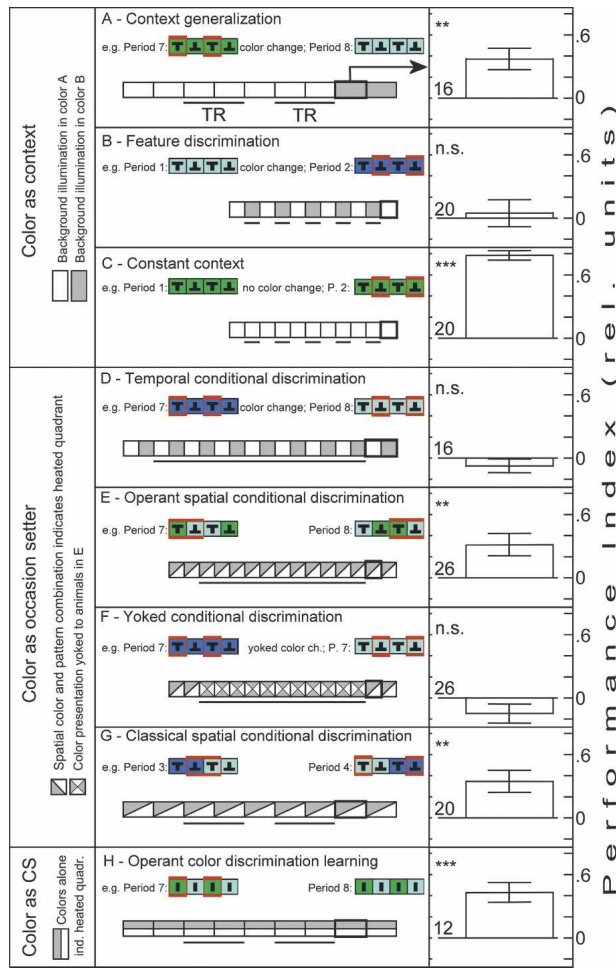


Figure 2. Colors can be context, occasion setters, and conditioned stimuli, depending on the temporal arrangement with the unconditioned stimulus. (A,B,C) Color as context. The color of background illumination during operant visual pattern discrimination learning changes according to the experimental schedule. (A) Context generalization. A single change of background illumination after the final training period marks the beginning of a 2-min test period for pattern memory in the new background color. The flies are able to show pattern memory in the new context. (B) The color changes are concomitant with the change from training to test periods. Increasing the number of context changes with respect to A abolishes the generalization effect—feature discrimination. (C) Control group in which training and test periods alternate in constant background color. Alternating training and test periods as in B does not abolish the memory score. (D–G) Color as occasion setter. (D) Color changes indicate the reversal of the pattern-heat contingency—temporal conditional discrimination. Reversal learning cannot be facilitated by context changes. (E,F,G) Colors change independently of the experimental schedule and indicate heated quadrants in conjunction with visual patterns. (E) Flies can solve a fully operant spatial conditional discrimination paradigm. (F) Flies fail to solve a temporal conditional discrimination task, where the color presentations are yoked to the animals in E. (G) Flies can solve spatial conditional discrimination with classical training and operant test periods. (H) Color as CS. The colors used in this study can be discriminated by wild-type flies in an operant visual learning task, with the colors as conditioned stimuli. Colored boxes with patterns illustrate the experimental design as in Figure 1 (see Materials and Methods for details). White/gray squares indicate 1-min. periods, rectangles 2-min. periods in the experimental time course. The performance indices of the highlighted test periods (bold) are displayed in the bar-graphs on the right. *** $P < 0.001$; ** $P < 0.01$; * $P < 0.05$; (n.s.) not significant. Numbers next to bar graphs indicate number of animals. Lines under experimental periods (indicated by “TR” in A) denote training periods. Performance index: $PI = (t_a - t_b)/(t_a + t_b)$.

avoid the heated patterns only in the illumination in which the patterns were actually combined with heat. Thus, the flies are able to detect that the change in coloration can predict a change in the heating regime (for more details see Wiener 2000). One could say that in this procedure, the colors that used to be part of the context are now “setting the occasion” for the punishment of certain pattern orientations, such that a pattern memory that at first was context-independent (Fig. 2A) becomes now context-dependent (Fig. 2B). In this interpretation, the paradigm is classified as feature discrimination, a case of occasion setting (Bouton and Nelson 1994; Holland et al. 1997; Dibbets et al. 2002). However, more experiments are needed to unambiguously conclude occasion setting as a mechanism for solving this feature discrimination task. We conclude that pattern memory can be both context independent and dependent, depending on the temporal relationship of the involved stimuli.

Colors in conditional discrimination tasks

Taking these experiments an additional step further, we trained the animals to avoid one set of patterns in one color (say the upright T's in green color) and the other set in another color (for instance, the inverted T's in blue-green), using the same alternating schedule. Such a paradigm is also a case of occasion setting and often referred to as conditional discrimination (Rescorla et al. 1985; Colwill et al. 1988a; Wilson and Pearce 1990). In our case, the colors change between experimental periods, rendering this experiment a case of temporal conditional discrimination. No significant learning scores were obtained, not even after 14 min of training (see Fig. 2D, Temporal conditional discrimination). Most learning effects at the *Drosophila* flight simulator become asymptotic after 8 min of training (Brembs and Heisenberg 2000). Of course, with a negative result it is impossible to rule out that the effect will appear with additional training (for discussion, see Brembs and Heisenberg 2001). We conclude that even the most amount of training used in this study was not able to reveal occasion setting when the CS was controlled operantly and the OS was presented in a temporal sequence (i.e., classically). Below we describe the results of an alternative training and testing regime of temporal conditional discrimination, which also did not yield a significant learning score (yoked conditional discrimination). Simultaneously, this experiment falsifies again the hypothesis that colors and patterns form unique percepts (compounds) in this case, one would expect a significant learning score as the paradigm corresponds to a simple discrimination.

In the previous experiment, only the patterns were under operant control, the colors were changed every minute by the computer program. Operant control of CSs facilitates learning about them (Brembs 2000; Brembs and Heisenberg 2000; Heisenberg et al. 2001). Possibly, operant control of the background colors would also lead to a facilitation of learning about their predictive value. Therefore, in a fourth step, we developed a scheme where the flies controlled both colors and patterns operantly. In this scheme, the color switches between quadrant borders such that flying toward an upright T will be punished in one color and unpunished in the other; the reverse contingency holds for the inverted T. In contrast to the previous experiment, the color changes take place at fixed points in space, which is why we termed it spatial conditional discrimination. The significant learning score indicates a successful operant spatial conditional discrimination experiment (Fig. 2E, Operant spatial conditional discrimination).

The standard experiment determining the importance of operant control is a so-called “yoked” control, where the stimuli

under investigation are delivered to the animal in the same temporal sequence in which operantly trained animals experienced them. In our version of the yoked control, the patterns were under operant control of the fly, but the colors were changed according to the sequence stored during the experiment (Fig. 2E), where both colors and patterns had been under operant control (a so-called replay experiment) (Wolf and Heisenberg 1991; Brembs and Heisenberg 2000). Whenever the color changed, the pattern-heat association was reversed according to the stored sequence of color changes by the previously trained flies. The animals could thus control their flight direction with respect to the patterns operantly, but the pattern-heat association changed with the illumination of the arena independently of their behavior (Fig. 2F, Yoked conditional discrimination). For instance, if the fly was flying toward an unpunished upright T in blue-green arena illumination, the color would change to blue (and the heat turn on) whenever the animal it was yoked to had changed arena coloration (without changing the position of the arena for the current fly). We favored this approach over a yoked control where both color and pattern sequences were played back, because it has already been shown that pattern replay does not support conditioning in the timeframe used here (Wolf and Heisenberg 1991; Brembs and Heisenberg 2000). Training with operant control of the patterns and yoked presentation of the colors did not yield significant learning scores (Fig. 2F). This result is in line with the failure to acquire a temporal conditional discrimination (Fig. 2D), as the yoked procedure also changes the color background according to a temporal regime, albeit not at regular intervals, but at the time points specified in the stored sequence. In this respect, the yoked control experiment amounts to a transfer experiment in which the flies are trained in a temporal conditional discrimination paradigm and are tested in a spatial conditional discrimination paradigm. Interestingly, the yoked control group also exhibits significantly lower heat avoidance (repeated measures ANOVA, SS: 81.44, df: 1, MS: 81.44, $F_{(278,4)}$, $P < 0.001$; data not shown). Apparently, exafferent inversion of the pattern-heat contingency is detrimental for heat avoidance, even if it is signaled by a concomitant change in background coloration. In conclusion, the yoked control experiment demonstrates that it is neither the number nor the dynamics of the color changes that lead to successful conditional discrimination when the animals control both colors and patterns operantly (Fig. 2E). Consequently, the arrangement of operant control over both patterns and colors enabled the animals to learn that color predicts the nature of the pattern-heat contingency.

Operant control facilitates conditional discrimination

We also investigated whether spatial conditional discrimination could be accomplished entirely classically, i.e., by training the animals with the same spatial arrangement of colors and patterns as in the fully operant experiment, but independently of their behavior. To this end, the arena was rotated slowly around the animal in open loop. Colors and heat were switched according to the same rules as during the fully operant experiment, i.e., between the patterns. Thus, the animals could learn about the colors predicting the nature of the pattern-heat contingency, as in the other conditional discrimination tasks, but this time the colors were arranged with the same spatial relationship to the arena position as in the fully operant experiment. *Drosophila* can, in principle, be classically conditioned to learn this occasion setting situation (Fig. 2G, Classical spatial conditional discrimination). However, the classical procedure was performed by exposing the flies to equal amounts of heat and no-heat, whereas an operant occasion setting yielded significant results with the flies avoiding the heat for about 86% of the training periods (average training $PI = 0.72$). Possibly, even more training would also lead to a sig-

nificant learning effect in the temporal conditional discrimination task (Fig. 2D). This difference in heat requirement indicates a more efficient conditional discrimination if all predictive stimuli are under operant control. Previous work on operant and classical discrimination learning has led to the hypothesis that operant control of environmental stimuli (i.e., composite operant conditioning) facilitates the acquisition of classical (i.e., CS-US) associations (Brembs 2000; Brembs and Heisenberg 2000; Heisenberg et al. 2001). Our results seem to suggest that learning about OSs is also facilitated by operant control.

There remains only one possible explanation of our conditional discrimination results thus far. The animals might detect which of the two color changes is associated with heat and no-heat, respectively, and then chose flight directions with respect to these references exactly between two patterns. In this view, the orientation of the patterns (i.e., upright or inverted T) would be irrelevant. We approached this possibility analytically and experimentally (Fig. 3). First, we plotted the time spent in each heated and nonheated semicircle of the arena, respectively. A fixation peak at the angular position exactly between two patterns (i.e., where the colors switch) would support the alternative explanation. However, the flies chose flight directions preferentially toward the patterns and not between two patterns (Fig. 3A). Nevertheless, the flies might only need a short period (or even only one instance) of switching from one color to the other in order to detect the safe flight directions, only to then continue to fixate the "safe" patterns. Therefore, we have conducted a control experiment identical to the operant spatial conditional discrimination paradigm (Fig. 2E) but with the patterns replaced by four identical, vertical stripes. If the flies learn to solve the operant spatial conditional discrimination paradigm by a simple association between turning direction, color switch, and heat, the pattern orientations should be irrelevant and the flies perform equally well in this control procedure. However, the flies do not show a significant learning score, falsifying this hypothesis and lending additional support to our hypothesis that indeed the logical combination between pattern orientations and color is learned in our case of occasion setting (Fig. 3B).

Colors as conditioned stimuli

Finally, in the fifth step, we conducted an operant color-discrimination learning experiment similar to the one developed earlier (Wolf and Heisenberg, 1997). In this experiment, the same colors that were used as context and OS in the experiments described above were set up to predict the occurrence of each single heating episode (i.e., the colors were set up as CS). In order to successfully solve this learning task, the fly has to choose "safe" flight directions with respect to four identical stripes. The only predictor of punishment is the arena coloration that switches between the identical stripes according to the conditioning schedule. *Drosophila* can learn to associate one of the colors with heat and avoid flight directions that lead to an arena illumination of this color (Fig. 2H, Operant color discrimination learning). A priori, this was not necessarily to be expected, considering that the colors used by Wolf and Heisenberg (1997), in contrast to the colors used here, do not support context generalization when used as context (Liu et al. 1999). However, in light of the successful conditional discrimination experiments, the suitability of these colors as CSs comes as somewhat less of a surprise. This result also emphasizes that the flies can readily distinguish the colors, even though they do not reveal that they can detect the color change in the wild-type context generalization experiment.

Summarizing the wild-type results, it emerges that the color pairs used for successful context generalization (i.e., blue-green/blue and blue-green/green) are also suitable to serve both as OS




Figure 3. Operant conditional discrimination learning is not due to simple conditioning. (A) Summed fixation histograms of the last 7 min for all 26 flies in the operant conditional discrimination experiment depicted in Figure 2E. (Left) Relative time spent at flight directions of 0–360° with respect to one of the hot/cold borders during the last 5 min of training. (Right) Flight directions during the following two minutes of test. The flies fixate the patterns preferentially both during training and during test. (Shaded area) Quadrants associated with heat; (white area) quadrants associated with heat-off; (dashed lines) quadrant borders between T-patterns. (B) Performance index of a control experiment in which patterns and colors were arranged as in an operant spatial conditional discrimination experiment (Fig. 2E), but with the T-shaped patterns replaced by four identical stripes as in color discrimination learning (Fig. 2H). No significant performance index was obtained. Performance index: $PI = (t_a - t_b) / (t_a + t_b)$.

and as CS. It can be concluded that a contextual stimulus can in principle become OS or CS, depending only on its temporal relationship to the US. In this light, both context dependence and independence appear as features of the brain.

Mushroom bodies are required for context generalization, but are dispensable for conditional discrimination

Knowing that flies without functional mushroom bodies fail to generalize between contexts (Fig. 4A, Context generalization; Liu et al. 1999), but do learn colors and patterns as CSs (Wolf et al. 1998), we tested transgenic flies with blocked mushroom-body output for their ability to perform conditional discrimination (Fig. 4B,C). Possibly, flies without functional mushroom bodies only fail to generalize between contexts because they learn quicker than wild-type flies that there will not be any punishment in the new context (see Figs. 2A, 4A; in other words, flies without mushroom bodies may exhibit facilitated conditional discrimination). The unsuccessful temporal conditional discrimination experiment for these flies (Fig. 4B, Temporal conditional discrimination) falsifies this hypothesis, as one would have expected the facilitated learning to improve performance also in this task over wild type. The negative results from both wild-type (Fig. 2D) and transgenic flies (Fig. 4B) in this experiment also corroborate the findings (this study, as well as data from Liu et al. 1999; Brembs and Heisenberg 2001; Tang and Guo 2001) that colors and patterns are processed separately and not as compounds (even in flies with impaired mushroom-body function!).

Interestingly, flies with blocked mushroom-body output did produce significant performance indices after training in our fully operant, spatial version of the occasion setting paradigm (Fig. 4C, Operant spatial conditional discrimination).

Discussion

Higher-order learning in *Drosophila*

For a learning situation to be classified as occasion setting, three criteria have to be met (Young et al. 2000; Pearce and Bouton 2001; Law et al. 2004). First, the two stimuli have to be processed individually and not treated as a compound. Second, the OS does not enter into an association with the US alone (no simple conditioning of the OS). Third, the OS has to be a specific modulator of the CS. We have corroborated previous evidence for the separate processing of colors and patterns (Brembs and Heisenberg 2001) by showing that our choice of colors and patterns are indeed separable (Fig. 2A). Indeed, we have shown that the animals have to learn to incorporate the colors into the pattern memory template (Fig. 2B). Additionally, one of our conditional discrimination paradigms (Fig. 2D, Temporal conditional discrimination) should have shown a significant learning score if colors and patterns had been processed as compounds. Moreover, even different pattern memories are processed by different layers of the fan-shaped body (Liu et al. 2006), making innate compound processing of pattern and color memories unlikely. Thus, there are several in-

dependent lines of evidence suggesting that the first criterion is met. In vertebrates, extinction or transfer experiments have conventionally been used to meet the second criterion (Bouton and Swartzentruber 1986; Holland 1986, 1989a,b; Myers and Gluck 1994; Young et al. 2000; Pearce and Bouton 2001; Law et al. 2004). These studies typically involved simple Pavlovian conditioning procedures with a single OS indicating the presence or absence of the US (i.e., feature discrimination, see, e.g., Bouton and Nelson 1994; Holland et al. 1997; Young et al. 2000; Pearce and Bouton 2001; Dibbets et al. 2002). These feature discrimination designs (e.g., Fig. 2B) do indeed require such additional controls, as the OS may be associated with the US alone. However, our final conditional discrimination designs are fully symmetrical (Fig. 2D,E,G), i.e., both OS and CS come in an equally non-predictive pair. Both the color pair and the pattern pair are by themselves equally associated with heat and therefore unsuitable as predictors of the heat—50% of both colors are associated with heat and 50% of both pattern orientations. Consequentially, neither of the two stimuli can enter into the association alone (simple conditioning cannot take place). Only by using the logical combination between the two (reminiscent of a configural learning task) (see, e.g., Pearce 1987, 1994; Young et al. 2000; Pearce and Bouton 2001), can the fly solve our conditional discrimination task (thus meeting the second criterion). Finally, our conditional discrimination designs use all possible permutations of colors and pattern orientations, excluding the possibility of the colors serving as a general (unspecific) modulator on the patterns—they must be learned as specific modulators of the pattern-heat contingency and thereby meet the third criterion.

We interpret these results as evidence that the flies solve the conditional discrimination tasks via an occasion-setting mechanism. To our knowledge, such higher-order learning has been found in only a few invertebrates (Colwill et al. 1988a; Rogers et




Figure 4. Flies with blocked mushroom-body output fail in context generalization and temporal conditional discrimination, but perform in spatial conditional discrimination. (A) Context generalization as in Figure 2A. Flies with blocked mushroom-body output do not transfer the pattern memory acquired during training to a different background color. (B) Temporal conditional discrimination where the colors in each period indicate the nature of the pattern/heat contingency (as in Fig. 2D). As in wild-type flies, color changes do not facilitate reversal learning and, hence, are not learned as occasion setters in flies with blocked mushroom-body output. (C) Operant spatial conditional discrimination as in Figure 2E, but with altered period set-up and duration. Colored boxes with patterns illustrate the experimental design as in Figure 1 (see Materials and Methods for details). White/gray rectangles indicate 2-min. periods, as in Figure 2. Performance indices of highlighted (bold) test periods are displayed in the bar-graphs on the right. $**P < 0.01$; (n.s.) not significant. Numbers next to bar graphs indicate number of animals. Lines under experimental periods (indicated by “TR” in A) denote training periods. Performance index: $PI = (t_a - t_b)/(t_a + t_b)$.

al. 1996; Schubert et al. 2002; Law et al. 2004; Matsumoto and Mizunami 2004) and has never been shown to be facilitated by operant control of the involved stimuli. This opens up higher-order conditioning for investigation in a genetically tractable, small-brained model system. Interestingly, free-flying honey bees are able to solve an operant conditional discrimination task (Schubert et al. 2002), whereas such a paradigm has proven more difficult to develop and is still lacking in the classical proboscis extension reflex with harnessed bees.

Mushroom bodies stabilize memory templates during context changes

Our approach opens the exciting aspect to investigate the neural mechanisms of higher-order associative processing. In simple associative conditioning, temporal coincidence is sufficient for individual neurons to mediate the learning process (Carew et al. 1981; Kandel and Schwartz 1982; Walters and Byrne 1983; Murphy and Glanzman 1999; Brembs et al. 2002; Antonov et al. 2003). In occasion setting, each stimulus by itself will trigger coincident firing where its pathway and the one of the US converge, since each separately will be reinforced—only the logical combination determines an unambiguous rule. Hence, the biological modifications in the brain need to take place in brain regions where the logical connection between OS and CS is preserved and not where the CS or the OS themselves are processed. Interestingly, memory traces for visual patterns have been localized to different layers of the fan-shaped body (Liu et al. 2006), while the brain regions implicated in the processing of CS and context are the mushroom bodies (Liu et al. 1999). Among a variety of measures that compromise mushroom-body function, we chose to use transgenic expression of tetanus toxin in mush-

room-body Kenyon cells to prevent synaptic transmission. Flies with compromised mushroom-body function perform well in a range of behaviors. They show coordinated walking, the full male courtship sequence, visual flight control, and basic responses to various stimuli (Heisenberg et al. 1985; Heisenberg and Wolf 1988; de-Belle and Heisenberg 1994; Connolly et al. 1996). While they can solve a number of learning tasks (Wolf et al. 1998), they fail in context generalization (Liu et al. 1999). In a first demonstration of the powerful combination of a closed behavioral methodology and modern transgenics, we have subjected flies with blocked mushroom-body output to both context generalization and occasion-setting experiments. While the flies failed to express pattern memory after a context change (Fig. 4A) (extending previous results from Liu et al. 1999), they could use the predictive information in the context change to learn about the pattern-heat contingency and solve the operant conditional discrimination task (Fig. 4C). Although we consider it unlikely, it would nevertheless be interesting to test more mushroom-body-specific driver lines to possibly find strains that fail in both tasks. Additionally, converging results from redundant techniques such as the ablation of the mushroom bodies by

treating larvae with hydroxy-urea (de-Belle and Heisenberg 1994) would serve to corroborate any findings in other transgenic lines. The mushroom bodies are a prominent neuropil and a hotspot of research (Heisenberg 2003; Gerber et al. 2004). Among many hypotheses, it has been proposed that mushroom bodies reduce the sensitivity to context changes by first extracting the CS from the context and then stabilizing the CS-US memory template against context changes (Liu et al. 1999). One original hypothesis was that flies without functional mushroom bodies cannot extract the CS from the context as well as wild-type flies and, hence, are not able to express the memory in the new context (Liu et al. 1999). An alternative hypothesis explains the context dependence in mushroom-body-impaired flies with enhanced occasion setting. Our temporal conditional discrimination task can discriminate between these alternatives. If mushroom bodies are involved in the separation of CS from context, mushroom-body-impaired flies should at least show a small learning score after 14 min of training (Fig. 4B), as they only have to solve two simple conditioning tasks instead of a higher-order task. If mushroom bodies reduce the capacity to learn occasion setting, mushroom-body-impaired flies should show a significant learning score already after <14 min of training. However, the mushroom-body-impaired flies completely failed this task, just as the wild-type flies. Suppose our results also hold for other transgenic lines as well as other, redundant techniques, what would this mean for our understanding of mushroom-body function?

Our data show that flies with impaired mushroom-body output probably can extract the patterns from the color background (Fig. 4B). If the colors were just part of the memory template, as Liu et al. (1999) suggest, the temporal conditional discrimination paradigm should amount to two simple conditioning tasks, e.g., the upright T in blue background is something

entirely different from the upright T in blue-green background. Flies with impaired mushroom-body function can solve such simple conditioning tasks (Wolf et al. 1998) and there should be no interference between the two—flies can store at least four memory templates simultaneously (Heisenberg et al. 2001). Thus, there appears no reason why these flies should not, in principle, be able to solve the task if patterns and colors were indeed not separable for them. It is tempting to even go so far as to interpret this result as evidence for context-independent memory in mushroom-body-impaired flies, but further experiments are needed to corroborate this interpretation.

On the same grounds, the data also rule out the alternative hypothesis that flies with impaired mushroom-body function are more efficient in occasion-setting tasks. If the reduction in avoidance after a context change (Fig. 4A) were indeed due to occasion setting, the temporal conditional discrimination task (Fig. 4B) should show a significant learning score already early in the experiment.

In conclusion, flies without mushroom bodies still appear to separate colors and patterns, but do not show enhanced occasion setting. With such evidence against both hypotheses, the picture instead emerges that mushroom bodies may specifically enhance the stability of memory traces against changes in the stimulus situation, but do not decrease the ability to detect such changes. Higher-order learning (occasion setting) appears to be independent of this function of the mushroom bodies. There is evidence suggesting that the reason for this independence lies in the different processing of generalization vs. discrimination tasks (Brembs and Hempel de Ibarra 2006).

Context dependence can be both feature and failure

One can intuitively understand the value both in heeding and in ignoring a signal for a biologically important event in a new situation. The former recognizes that the signal may still be relevant, even in the new situation; the latter saves valuable resources by recognizing that in the new situation the signal needs to validate its signaling qualities. One can also easily understand that the degree to which the two situations differ has an effect on the outcome of such experiments. Indeed, the choice of colors is important for context-independent or context-dependent pattern memory in *Drosophila* (Liu et al. 1999). Presumably, and this is precisely where the issue becomes more intricate, the crucial determinant of context dependence or independence are the (species-specific) stimuli that make up what the human observers call “context.” For some researchers, context generalization is the brain capacity and biological manipulations are sought to compromise it (e.g., Liu et al. 1999). For others, the context dependence is the brain capacity (occasion setting) and biological manipulations are sought to compromise it (e.g., Law et al. 2004). The conundrum is not enlightened by a large number of terms denoting the same or very similar experimental situations. Context generalization (this study; Liu et al. 1999; Brembs and Hempel de Ibarra 2006) takes place when a memory can be transferred between different contexts (context-independent memory). Contextual learning or context conditioning (e.g., Colwill et al. 1988b; Kim and Fanselow 1992; Rogers and Matzel 1996; Debiec et al. 2002) refers to experiments in which different contexts predict the occurrence or nonoccurrence of otherwise unsigned USs. Feature discrimination (e.g., Bouton and Nelson 1994; Holland et al. 1997; Dibbets et al. 2002) refers to experiments where a stimulus (or context) predicts the reinforcement/nonreinforcement of a CS. Conditional discrimination (e.g., Rescorla et al. 1985; Colwill et al. 1988a; Wilson and Pearce 1990), trans-switching (e.g., Furedy 1991; Lachnit and Kimmel 1991), and ambiguous discrimination (e.g., Holland 1991) all denote

experiments that are also classified as occasion setting (e.g., Swartzentruber 1991; Bonardi and Hall 1994; Miller and Oberling 1998; Schmajuk et al. 1998; Young et al. 2000; Clarke et al. 2001)—a stimulus (or context) characterizes the nature of a CS-US contingency or discrimination. Traditionally, researchers distinguished between predictive stimuli and mere “context” either by the physical properties of the stimuli (e.g., Bouton et al. 1999), or according to their temporal relationship to the US (e.g., Wickens 1987). The work on *Drosophila* at the flight simulator demonstrated that both physical properties (Liu et al. 1999; Brembs and Hempel de Ibarra 2006) and the nature of the predictive relation to the reinforcer (this study) are critical for the decision of whether to treat two situations as equivalent or as fundamentally different. The fewer the changes between situations, the more pronounced the impact of the physical properties of the situation (Liu et al. 1999; Brembs and Hempel de Ibarra 2006); the more changes, the more pronounced the role of the changes and their relationship to the reinforcement (this work; see also Swartzentruber 1991; Myers and Gluck 1994). Thus, for the general organization of learning experiments, it must be emphasized that the classification of stimuli as “context” is less obvious and self-explanatory than it might seem. Moreover, whether the non-retrieval of a memory in any however slightly altered experimental situation can be considered a feature (contextual memory) or a failure (no generalization) cannot be addressed without further experiments of the kind detailed in this work. The overarching brain capacity is to be able to flexibly generalize or discriminate between two situations depending on the information the difference between the situations conveys to the animal. Investigating generalization or discrimination individually is a one-sided endeavor and may thus yield confusing results. In our companion paper, we have applied this new methodology and addressed the inter-dependence of the physical parameters and the predictive value to show that discrimination and generalization of background colors are supported by different parameters in *Drosophila* (Brembs and Hempel de Ibarra 2006).

For fruit flies, as for humans, the claim appears valid that “like parallel research on occasion setting, research on contextual control suggests that a more complex associative structure may often be acquired in associative learning.” (Pearce and Bouton 2001).

Materials and Methods

Flies

Flies are kept on standard cornmeal/molasses medium (Guo et al. 1996) at 25°C and 60% humidity with a 14-h light/10-h dark regime. Females aged 24–48 h are briefly immobilized by cold-anesthesia and glued (Loctite UV glass glue) with head and thorax to a triangle-shaped copper hook (diameter 0.05 mm) the day before the experiment. The animals are then kept individually overnight in small moist chambers containing a few grains of sucrose.

Transgenes

Sweeney et al. (1995) developed a method that constitutively blocks synaptic transmission by expressing the catalytic subunit of bacterial tetanus toxin (Cnt-E) in target neurons in the *Drosophila* brain using the P[GAL4] technique (Brand and Perrimon 1993). Because of the effects of mushroom-body function on context generalization (Liu et al. 1999), we use the Cnt-E transgene to block synaptic output from the mushroom bodies. The P[GAL4] line mb247 (Schulz et al. 1996) is used as a mushroom-body-specific GAL4 driver (Zars et al. 2000). This driver strain has not been tested for context generalization previously. We use the trans-heterozygote offspring from the driver (mb247) and the

reporter strain (UAS_{GAL4}-Cnt-E) for our studies as described previously (Sweeney et al. 1995; Baier et al. 2002).

Apparatus

The *Drosophila* flight simulator is a computer-controlled feedback system in which the fly uses its yaw torque to control the rotations of a panorama surrounding it (Fig. 1). The core device is the torque meter (Götz 1964; Heisenberg and Wolf 1984), which measures a fly's angular momentum around its vertical body axis. The fly, glued to the hook, is attached to the torque meter via a clamp to accomplish stationary flight in the center of a cylindrical panorama (arena; diameter 58 mm), homogeneously illuminated from behind (Fig. 1). The light source is a 100W, 12V tungsten-iodine bulb. For background coloration of the arena, the light is passed through one of three different broad band filters—(1) broadband blue (Kodak Wratten gelatin filter No. 47); (2) broadband green (Kodak Wratten gelatin filter No. 99); and (3) “daylight” blue-green (Rosco “surfblue” No. 5433). The transmission spectrum of the Rosco blue-green filter used in this study is equivalent to that of the BG18 filter (Schott, Mainz) used by Liu et al. (1999) (data not shown). Filters can be exchanged by a fast solenoid within 0.1 sec.

A computer-controlled electric motor rotates the arena such that its angular velocity is proportional to, but directed against the fly's yaw torque (coupling factor $K = -11^\circ/\text{sec} \cdot 10^{-10} \text{ Nm}$). This enables the fly to stabilize the panorama and to control its angular orientation. This virtual “flight direction” (i.e., arena position) is recorded continuously via a circular potentiometer (Novotechnik, A4102a306). An analog to digital converter card (PCL812; Advantech Co.) feeds arena position and yaw torque into a computer that stores the traces (sampling frequency 20 Hz) for later analysis.

Punishment is achieved by applying heat from an adjustable infrared laser (825 nm, 150 mW), directed from behind and above onto the fly's head and thorax. The laser beam is pulsed (~200 msec pulse width at ~4 Hz) and its intensity reduced to assure the survival of the fly.

General experimental design

Each fly is used only once. The time-course of the experiment is divided into consecutive periods of either 1 or 2 min duration. Depending on whether heat may be applied during such a period, it is termed a training period (heating possible) or a test period (heat off). Note that this nomenclature is independent of any higher-order training that may encompass several training/test periods. The treatment of the flies during these periods determines the type of experiment, as described below. Color pairs were always green/blue-green and blue/blue-green (Fig. 1).

Discrimination learning—patterns

For patterns as CS (Wolf and Heisenberg 1991), four black, T-shaped patterns of alternating orientation (i.e., two upright and two inverted) are evenly spaced on the arena wall (pattern width $\psi = 40^\circ$, height $\theta = 40^\circ$, width of bars = 14° , as seen from the position of the fly). A computer program divides the 360° of the arena into four virtual 90° quadrants, the centers of which are denoted by the patterns. During training periods, heat punishment is made contiguous with the appearance of one of the pattern orientations in the frontal visual field. Reinforcement of each pattern is always equalized within groups. During test periods, the heat is permanently switched off.

Discrimination learning—colors

For colors as CS (Wolf and Heisenberg 1997) the centers of the four virtual quadrants are denoted by four identical vertical stripes (width $\psi = 14^\circ$, height $\theta = 40^\circ$). The color of the illumination of the whole arena is changed whenever one of the virtual quadrant borders passes a point in front of the fly. During training periods, heat punishment is made contiguous with one color of the pairs blue/blue-green and green/blue-green. Reinforcement of each color is always equalized within groups. During test

periods, the heat is permanently switched off. See Figures 1 and 2H, Operant color discrimination learning.

Context generalization

Testing for the stability of pattern memory, the number of background color changes with each training/test period is varied.

1. Following the original context generalization experiment by Liu et al. (1999), only one color change takes place after seven 2-min periods ($2 \times$ test, $2 \times$ training, test, $2 \times$ training), before the final 2-min test period. The color is changed either between green and blue-green for half of the cases or between blue and blue-green for the rest of the cases, such that each color is training or test color in 25% of all experiments. A successful generalization experiment is characterized by a positive learning score, which indicates that the pattern memory was generalized across the different color contexts. Such a successful experiment also shows that the pattern can be processed independently from the color and the two stimuli are not perceived as a compound (Brembs and Heisenberg 2001). Context generalization is different from context conditioning where the animals learn to respond to a context. In this study, we never performed context conditioning, but only tested for the ability of a context change to disrupt the transfer of operant pattern memory between contexts. Successful context generalization is characterized by a lack of response to the context change. See Figures 1, 2A, and 3, Context generalization.
2. In a modification of the context generalization experiment, the number of context changes is increased from one (before the final test period; see above) to 10 (between every training and test period; see below). Specifically, the duration of experimental periods is reduced to 1 min and training and test periods alternated. In this manner, five training and five test phases alternate after 1 min pre-test in the test color. Such a design increases the predictive value of the colors, as they indicate whether flying toward the pattern is punished or not. That is, if the arena is illuminated with one color, flying toward one of the patterns is punished; if illuminated with the other color, none of the patterns is punished. The use of the pattern pairs (blue/blue-green and green/blue-green) is balanced between animals. After training in this manner, the animals are subsequently tested for pattern preference in both colors, with the heat permanently switched off. The sequence of training and test coloration is balanced between animals. See Figures 1 and 2B, Context indicates heat on/off. In a control procedure, animals were subjected to pattern discrimination learning with stationary colors (Figures 1 and 2C, Constant context).

Occasion setting

In a modification of the previous experiment, arena coloration is used to indicate the nature of the pattern-heat contingency. For instance, flying toward the upright T is punished under green illumination and the inverted T is unpunished, but then blue-green illumination indicates the reverse pattern-heat contingency. In this experiment, neither of the stimuli alone can unambiguously predict reinforcement. The animals can master this discrimination only by considering the different combinations of the stimuli. We use four different occasion setting paradigms as described in the following paragraphs.

1. Temporal conditional discrimination

This version is the direct extension of the context generalization experiments described above. By alternating 1-min periods, the patterns remain under operant control. After two test periods with each arena illumination color, 14 periods of training follow. The arena illumination color changes with each period. Pattern-heat contingencies alternate with the colors (colors balanced) and periods. In this manner, the patterns stay under operant control of the fly for the duration of the experiment, but the colors are under the control of the experimenter and therefore

presented classically, i.e., independently of the behavior of the animal. In the final two 1-min periods, the animals are tested for their pattern preference in the two colors. The PIs of the two final test periods are averaged and tested against zero. See Figures 1, 2D, and 3B, Temporal conditional discrimination.

2. Operant spatial conditional discrimination

In this paradigm the colors do not change between periods (temporally) but between quadrants (spatially). In this way, both colors and patterns come under operant control. The center of each quadrant is still denoted by the patterns (alternating upright and inverted Ts). The difference to temporal conditional discrimination consists of the arrangement of color and heat with the quadrants. While heat was associated with two opposite quadrants (e.g., the ones with the upright T in the center) before, heat is now associated with adjacent quadrants, i.e., one with an upright and one with an inverted T. Thus, instead of being switched on or off at each of the four quadrant borders, the heat is now switched on or off at only two opposite borders. The color of the arena illumination is changed at the remaining two quadrant borders, where the heat is not switched on or off. Thus, heat is applied in two quadrants, which include an upright and an inverted T as well as the quadrant border where the background coloration is changed. Conversely, arena coloration is changed exactly between the two punished patterns and between the two unpunished patterns. In such a way, the flies get heated when they fly toward, say, a green upright T and a blue-green inverted T and switch the heat off by flying into one of the other two quadrants with a green inverted T and a blue-green upright T. One arrangement of quadrants may thus look as follows: The first quadrant features the upright T and whenever the fly enters this quadrant, the whole arena turns to blue illumination. The second quadrant features the inverted T and the arena illumination remains blue. If the fly enters the third quadrant with the upright T, the whole arena turns to blue-green. In the fourth quadrant, the inverted T is in the center, but the arena illumination stays blue-green. The heat regime is such that neither pattern nor color alone could predict reinforcement. For example, heat is switched on whenever the fly enters quadrants 2 or 3, but no heat is presented when entering quadrants 1 or 4. This heat regime is used for half of the animals, whereas the other half of the animals is not punished in quadrants 2 or 3, but quadrants 1 or 4 are punished.

The training phase lasts 11 min, and is divided into 1-min periods. After each period, the arena is set to a random position to minimize conditioning to spurious spatial cues. The spatial arrangement of patterns and colors was randomized across periods (i.e., if the patterns in quadrant 1 and 2 were "blue" and the patterns in quadrant 3 and 4 "blue-green" in one period, this association was reversed in a random selection of other periods). This randomization minimized the spatial contingency and emphasized the logical contingency between patterns, heat, and colors. After 11 min of training, the animals are tested for 1 min for their quadrant preference with the heat permanently switched off. See Figures 1 and 2E, Operant spatial conditional discrimination. Transgenic flies (Fig. 4B, Operant spatial conditional discrimination) were trained for 8 min (four 2-min training periods) and tested for 2 min in the same temporal order as described for context generalization (Liu et al. 1999), "Classical spatial conditional discrimination" (see below) and color discrimination learning (see above). In a control experiment (Fig. 3B), the four T-patterns were replaced by four identical, vertical stripes as in color discrimination learning (see above).

3. Yoked conditional discrimination

This experiment aims to test for the requirement of operant control of the colors. The animals control the position of the patterns, but the sequence of color changes is played back from the animals previously trained in the previous experiment (patterns and colors operant), where they controlled both the patterns and the colors operantly. Thus, the dynamics and frequency of color changes are identical to the one where arena coloration is con-

trolled by the fly. Only in this experiment, the color change is independent of the animal's behavior. Every change in the background coloration reverses the pattern-heat contingency as before. This experiment is almost identical to the temporal conditional discrimination experiment above, with the exception of the temporal sequence of color changes matching those of spatial conditional discrimination and the test being identical to that of spatial conditional discrimination. After 11 min of training, the animals are tested for 1 min for their pattern-color preference with patterns and colors under operant control, as in the previous experiment (patterns and colors operant). See Figures 1 and 2F, Yoked conditional discrimination.

4. Classical spatial conditional discrimination

This experiment is similar to the one where patterns and colors are under operant control (patterns and colors operant). However, the arena is not under the fly's control, but is instead rotated at 30°/sec, such that the full 360° of patterns, accompanying color changes, and heat are experienced by the fly, but independently of the fly's behavior. By these means it has previously been shown that flies cannot only learn to discriminate patterns (Brembs and Heisenberg 2000) but also colors (data not shown) independently of their behavior (i.e., classically). Thus, this experiment differs from the temporal conditional discrimination in both the number of color changes as well as the spatial relationship of the color change and the arena position. It has been shown previously that flies are sensitive to such spatial information (Brembs and Heisenberg 2000). As in the fully operant experiment, the spatial arrangement of patterns and colors was randomized across periods (i.e., if the patterns in quadrant 1 and 2 were "blue" and the patterns in quadrant 3 and 4 "blue-green" in one period, this association was reversed in a random selection of other periods). Again, this randomization is intended to minimize the spatial contingency and emphasize the logical contingency between patterns, heat, and colors. After 8 min of training, the animals are tested for 2 min for their quadrant preference with the heat permanently switched off. In this test, both colors and patterns are under full operant control, as in the previous two experiments. See Figures 1 and 2G, Classical spatial conditional discrimination.

Data evaluation and statistics

The color and/or pattern preference of individual flies is calculated as the performance index $PI = (t_a - t_b)/(t_a + t_b)$. During training periods, t_b indicates the time the fly is exposed to the heat and t_a the time without heat. During test periods, t_a and t_b refer to the times when the fly choose the formerly (or subsequently) unpunished or punished situation, respectively. Thus, a PI of 1 means the fly spent the entire period in the quadrants not associated with heat, whereas a PI of -1 indicates that the fly spent the entire period in the quadrants associated with heat. Accordingly, a PI of zero indicates that the fly distributed the time evenly between heated and nonheated quadrants. PIs from test periods are called "test PIs" or "learning scores." Learning scores were tested for significance using a *t*-test for single means against zero, following Liu et al. (1999).

Acknowledgments

Our special gratitude goes to Bertram Gerber for vital intellectual input not only into the development of the spatial conditional discrimination paradigm, but also into the writing of this manuscript. We are indebted to Gérard Lebouille, Syed Abid Hussaini, and our Monday journal club for commenting on an earlier version of the manuscript.

References

- Antonov, I., Antonova, I., Kandel, E.R., and Hawkins, R.D. 2003. Activity-dependent presynaptic facilitation and hebbian ltp are both required and interact during classical conditioning in *Aplysia*. *Neuron* 37: 135–147.
- Baier, A., Wittek, B., and Brembs, B. 2002. *Drosophila* as a new model organism for the neurobiology of aggression? *J. Exp. Biol.*

- 205:** 1233–1240.
- Bonardi, C. and Hall, G. 1994. Occasion-setting training renders stimuli more similar: Acquired equivalence between the targets of feature-positive discriminations. *Q. J. Exp. Psychol. B* **47**: 63–81.
- Bouton, M.E. and Nelson, J.B. 1994. Context-specificity of target versus feature inhibition in a feature-negative discrimination. *J. Exp. Psychol. Anim. Behav. Process.* **20**: 51–65.
- Bouton, M.E. and Swartztruber, D. 1986. Analysis of the associative and occasion-setting properties of contexts participating in a Pavlovian discrimination. *J. Exp. Psychol. Anim. Behav. Process.* **12**: 333–350.
- Bouton, M.E., Nelson, J.B., and Rosas, J.M. 1999. Stimulus generalization, context change, and forgetting. *Psychol. Bull.* **125**: 171–186.
- Brand, A.H. and Perrimon, N. 1993. Targeted gene expression as a means of altering cell fates and generating dominant phenotypes. *Development* **118**: 401–415.
- Brembs, B. 2000. An analysis of associative conditioning in *Drosophila* at the flight simulator. Ph.D. thesis, University of Würzburg, Würzburg, Germany.
- Brembs, B. and Heisenberg, M. 2000. The operant and the classical in conditioned orientation in *Drosophila melanogaster* at the flight simulator. *Learn. Mem.* **7**: 104–115.
- . 2001. Conditioning with compound stimuli in *Drosophila melanogaster* in the flight simulator. *J. Exp. Biol.* **204**: 2849–2859.
- Brembs, B. and Hempel de Ibarra, N. 2006. Different parameters support generalization and discrimination learning in *Drosophila* at the flight simulator. *Learn. Mem.* (this issue).
- Brembs, B., Lorenzetti, F.D., Reyes, F.D., Baxter, D.A., and Byrne, J.H. 2002. Operant reward learning in *Aplysia*: Neuronal correlates and mechanisms. *Science* **296**: 1706–1709.
- Carew, T.J., Walters, E.T., and Kandel, E.R. 1981. Classical conditioning in a simple withdrawal reflex in *Aplysia californica*. *J. Neurosci.* **1**: 1426–1437.
- Clarke, H.A., Skinner, D.M., and van der Kooy, D. 2001. Combined hippocampal and amygdala lesions block learning of a response-independent form of occasion setting. *Behav. Neurosci.* **115**: 341–357.
- Colwill, R.M., Absher, R.A., and Roberts, M.L. 1988a. Conditional discrimination learning in *Aplysia californica*. *J. Neurosci.* **8**: 4440–4444.
- . 1988b. Context-us learning in *Aplysia californica*. *J. Neurosci.* **8**: 4434–4439.
- Connolly, J.B., Roberts, I.J., Armstrong, J.D., Kaiser, K., Forte, M., Tully, T., and O’Kane, C.J. 1996. Associative learning disrupted by impaired *gs* signaling in *Drosophila* mushroom bodies. *Science* **274**: 2104–2107.
- de-Belle, J.S. and Heisenberg, M. 1994. Associative odor learning in *Drosophila* abolished by chemical ablation of mushroom bodies. *Science* **263**: 692–695.
- Debiec, J., LeDoux, J.E., and Nader, K. 2002. Cellular and systems reconsolidation in the hippocampus. *Neuron* **36**: 527–538.
- Dibbets, P., Maes, J.H., and Vossen, J.M. 2002. Serial and simultaneous feature positive discriminations in a human skinner box. *Behav. Processes* **59**: 1–14.
- Dickinson, A. and Balleine, B.W. 2002. The role of learning in the operation of motivational systems. In: *Learning, motivation and emotion, volume 3 of Steven’s handbook of experimental psychology*, 3rd ed., (ed. C.R. Gallistel), pp. 497–533. John Wiley & Sons, New York.
- Ernst, R. and Heisenberg, M. 1999. The memory template in *Drosophila* pattern vision at the flight simulator. *Vision Res.* **39**: 3920–3933.
- Furedy, J.J. 1991. On testing current cognitive theories of Pavlovian conditioning in the human Pavlovian autonomic transswitching preparation. *Biol. Psychol.* **32**: 201–204.
- Gerber, B., Tanimoto, H., and Heisenberg, M. 2004. An engram found? Evaluating the evidence from fruit flies. *Curr. Opin. Neurobiol.* **14**: 737–744.
- Götz, K.G. 1964. Optomotorische untersuchung des visuellen systems einiger augenmutanten der fruchtfliege *Drosophila*. *Kybernetik* **2**: 77–92.
- Guo, A., Liu, L., Xia, S.-Z., Feng, C.-H., Wolf, R., and Heisenberg, M. 1996. Conditioned visual flight orientation in *Drosophila*; dependence on age, practice and diet. *Learn. Mem.* **3**: 49–59.
- Heisenberg, M. 2003. Mushroom body memoir: From maps to models. *Nat. Rev. Neurosci.* **4**: 266–275.
- Heisenberg, M. and Wolf, R. 1984. *Vision in Drosophila*. *Genetics of microbehavior*. Springer, Berlin, Heidelberg, New York, Tokyo.
- . 1988. Reafferent control of optomotor yaw torque in *Drosophila melanogaster*. *J. Comp. Physiol. [A]* **163**: 373–388.
- Heisenberg, M., Borst, A., Wagner, S., and Byers, D. 1985. *Drosophila* mushroom body mutants are deficient in olfactory learning. *J. Neurogenet.* **2**: 1–30.
- Heisenberg, M., Wolf, R., and Brembs, B. 2001. Flexibility in a single behavioral variable of *Drosophila*. *Learn. Mem.* **8**: 1–10.
- Holland, P.C. 1986. Temporal determinants of occasion setting in feature-positive discriminations. *Anim. Learn. Behav.* **14**: 111–120.
- . 1989a. Acquisition and transfer of conditional discrimination performance. *J. Exp. Psychol. Anim. Behav. Process.* **15**: 154–165.
- . 1989b. Feature extinction enhances transfer of occasion setting. *Anim. Learn. Behav.* **17**: 269–279.
- . 1991. Transfer of control in ambiguous discriminations. *J. Exp. Psychol. Anim. Behav. Process.* **17**: 231–248.
- Holland, P.C., Hamlin, P.A., and Parsons, J.P. 1997. Temporal specificity in serial feature-positive discrimination learning. *J. Exp. Psychol. Anim. Behav. Process.* **23**: 95–109.
- Kandel, E.R. and Schwartz, J.H. 1982. Molecular biology of learning: Modulation of transmitter release. *Science* **218**: 433–443.
- Katsov, A. and Clandinin, T.R. 2006. Insect vision: Remembering the shape of things. *Curr. Biol.* **16**: R369–R371.
- Kim, J.J. and Fanselow, M.S. 1992. Modality-specific retrograde-amnesia of fear. *Science* **256**: 675–677.
- Lachnit, H. and Kimmel, H.D. 1991. Extending the Rescorla-Wagner theory to account for transswitching. *Biol. Psychol.* **32**: 193–200.
- Law, E., Nuttley, W.M., and van der Kooy, D. 2004. Contextual taste cues modulate olfactory learning in *C. elegans* by an occasion-setting mechanism. *Curr. Biol.* **14**: 1303–1308.
- Liu, L., Wolf, R., Ernst, R., and Heisenberg, M. 1999. Context generalization in *Drosophila* visual learning requires the mushroom bodies. *Nature* **400**: 753–756.
- Liu, G., Seiler, H., Wen, A., Zars, T., Ito, K., Wolf, R., Heisenberg, M., and Liu, L. 2006. Distinct memory traces for two visual features in the *Drosophila* brain. *Nature* **439**: 551–556.
- Matsumoto, Y. and Mizunami, M. 2004. Context-dependent olfactory learning in an insect. *Learn. Mem.* **11**: 288–293.
- Miller, R.R. and Oberling, P. 1998. Analogies between occasion setting and Pavlovian conditioning. In *Occasion setting: Associative learning and cognition in animals* (eds. N.A. Schmajuk and P.C. Holland), pp. 3–35. American Psychological Association, Washington, D.C.
- Murphy, G.G. and Glanzman, D.L. 1999. Cellular analog of differential classical conditioning in *Aplysia*: Disruption by the nmda receptor antagonist dl-2-amino-5-phosphonovalerate. *J. Neurosci.* **19**: 10595–10602.
- Myers, C.E. and Gluck, M.A. 1994. Context, conditioning, and hippocampal rerepresentation in animal learning. *Behav. Neurosci.* **108**: 835–847.
- Pearce, J.M. 1987. A model for stimulus generalization in Pavlovian conditioning. *Psychol. Rev.* **94**: 61–73.
- . 1994. Similarity and discrimination: A selective review and a connectionist model. *Psychol. Rev.* **101**: 587–607.
- Pearce, J.M. and Bouton, M.E. 2001. Theories of associative learning in animals. *Annu. Rev. Psychol.* **52**: 111–139.
- Rescorla, R.A., Grau, J.W., and Durlach, P.J. 1985. Analysis of the unique cue in configural discriminations. *J. Exp. Psychol. Anim. Behav. Process.* **11**: 356–366.
- Rogers, R.F. and Matzel, L.D. 1996. Higher-order associative processing in *hermissenda* suggests multiple sites of neuronal modulation. *Learn. Mem.* **2**: 279–298.
- Rogers, R.F., Schiller, K.M., and Matzel, L.D. 1996. Chemosensory-based contextual conditioning in *hermissenda crassicornis*. *Anim. Learn. Behav.* **24**: 28–37.
- Schmajuk, N.A., Lamoureaux, J.A., and Holland, P.C. 1998. Occasion setting: A neural network approach. *Psychol. Rev.* **105**: 3–32.
- Schubert, M., Lachnit, H., Francucci, S., and Giurfa, M. 2002. Nonelemental visual learning in honeybees. *Anim. Behav.* **64**: 175–184.
- Schulz, R.A., Chromey, C., Lu, M.F., Zhao, B., and Olson, E.N. 1996. Expression of the *d-mef2* transcription in the *Drosophila* brain suggests a role in neuronal cell differentiation. *Oncogene* **12**: 1827–1831.
- Sutton, R.S. and Barto, A.G. 1998. *Reinforcement learning*. MIT Press, Cambridge, MA.
- Swartztruber, D. 1991. Blocking between occasion setters and contextual stimuli. *J. Exp. Psychol. Anim. Behav. Process.* **17**: 163–173.
- Sweeney, S.T., Broadie, K., Keane, J., Niemann, H., and O’Kane, C.J. 1995. Targeted expression of tetanus toxin light chain in *Drosophila* specifically eliminates synaptic transmission and causes behavioral defects. *Neuron* **14**: 341–351.
- Tang, S. and Guo, A. 2001. Choice behavior of *Drosophila* facing contradictory visual cues. *Science* **294**: 1543–1547.
- Tang, S.M., Wolf, R., Xu, S.P., and Heisenberg, M. 2004. Visual pattern recognition in *Drosophila* is invariant for retinal position. *Science* **305**: 1020–1022.
- Walters, E.T. and Byrne, J.H. 1983. Associative conditioning of single sensory neurons suggests a cellular mechanism for learning. *Science*

- 219:** 405–408.
- Wickens, D.D. 1987. The dual meanings of context: Implications for research, theory and applications. In: *Memory and learning: The Ebbinghaus centennial conference* (eds. D.S. Gorfein and R.R. Hoffman) pp. 135–152. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- Wiener, J. 2000. Kontext-generalisierung in *Drosophila melanogaster*. Diploma thesis, Julius-Maximilians Universität, Würzburg, Germany.
- Wilson, P.N. and Pearce, J.M. 1990. Selective transfer of responding in conditional discriminations. *Q. J. Exp. Psychol. B* **42:** 41–58.
- Wolf, R. and Heisenberg, M. 1991. Basic organization of operant behavior as revealed in *Drosophila* flight orientation. *J. Comp. Physiol. [A]* **169:** 699–705.
- . 1997. Visual space from visual motion: Turn integration in tethered flying *Drosophila*. *Learn. Mem.* **4:** 318–327.
- Wolf, R., Wittig, T., Liu, L., Wustmann, G., Eyding, D., and Heisenberg, M. 1998. *Drosophila* mushroom bodies are dispensable for visual, tactile and motor learning. *Learn. Mem.* **5:** 166–178.
- Young, M.E., Johnson, J.L., and Wasserman, E.A. 2000. Serial causation: Occasion setting in a causal induction task. *Mem. Cognit.* **28:** 1213–1230.
- Zars, T., Fischer, M., Schulz, R., and Heisenberg, M. 2000. Localization of a short-term memory in *Drosophila*. *Science* **288:** 672–675.

Received May 14, 2006; accepted in revised form July 13, 2006.