

Flexibility in a Single Behavioral Variable of *Drosophila*

Martin Heisenberg,¹ Reinhard Wolf, and Björn Brembs²

Lehrstuhl für Genetik und Neurobiologie, Biozentrum, D-97074 Würzburg, Germany

The flexibility of behavior is so rich, and its components are so exquisitely interwoven, that one may be well advised to turn to an isolated behavioral module for study. Gill withdrawal in *Aplysia*, the proboscis extension reflex in the honeybee, and lid closure in mammals are such examples. We have chosen yawing, a single component of flight orientation in *Drosophila melanogaster*, for this approach. A specialty of this preparation is that the behavioral output can be reduced beyond the single module by one further step. It can be studied in tethered animals in which all turns are blocked while the differentially beating wings still provide the momentum. These intended yaw turns are measured by a torque meter to which the fly is hooked. The fly is held horizontally as if cruising at high speed. The head is glued to the thorax. It can bend its abdomen, extend its proboscis, and move its legs but cannot shift its direction of gaze or its orientation in space. Evidently, a fly hardly ever encounters this bizarre situation in the wild. We describe here the flexibility in this single behavioral variable. It provides insights into the relation between classical and operant conditioning, the processing of and interactions between the conditioned visual stimuli, early visual memory, visual pattern recognition, selective attention, and several other experience-dependent properties of visual orientation behavior. We start with a brief summary of visual flight control at the torque meter.

Yaw Torque in Tethered Flight

The torque meter (Götz 1964) provides a faithful on-line voltage analog of the fly's momentum around its vertical body axis. On inspecting these time traces, one notices that the fly does not keep a preferred value but, rather, explores the entire yaw torque range spontaneously—even in a featureless panorama. In addition to slow torque fluctuations, the fly generates short pulses of torque (torque spikes), the equivalent of the so-called body saccades observed in free flight. The slow torque fluctuations and torque spikes are generated in an apparently random temporal sequence, exemplifying the endogenous initiation of behavioral activity without any change in the fly's stimulus situation. A comparison of the time course of body saccades and torque spikes shows that tethering, indeed, has its price. Body saccades in free flight are much faster than those of tethered animals. This is because of proprioceptive feedback, probably from the halteres, monitoring (in free flight) the turning of the body and the cutting short of the torque pulses (Heisenberg and Wolf 1984).

While tethering truncates flight maneuvers, the torque meter allows the experimenter to manipulate their consequences by providing artificial sensory feedback. For in-

stance, one can divide the fly's spontaneous yaw torque range as recorded in the computer into two equally sized domains, roughly corresponding to left and right turns in free flight. One can now heat the fly whenever its yaw torque is in, for example, the right domain. In a matter of seconds, the fly learns to avoid this domain in order to avoid the heat and, after a few minutes, shows a lasting avoidance of that domain even if heat is switched off permanently (yaw torque [yt] learning; Wolf and Heisenberg 1991). This is a remarkable feat, as there are no other sensory stimuli contingent on punishment. Moreover, in no natural situation would continuous turning to one side bring the fly permanently out of a heated zone.

In a less simple but somewhat more realistic arrangement, one may simulate the visual movements that the fly's turning maneuvers would cause in free flight. The angular velocity the fly would acquire in free flight can be calculated from the on-line torque signal. But instead of turning the fly, the yaw torque is made to turn an arena around the fly in the opposite direction. This arrangement enables the fly to stabilize the arena, that is, to fly straight with respect to the patterns on the arena wall and to choose flight directions with respect to these patterns. The setup is called a flight simulator—a euphemism, because free flight resembles this situation very little.

The Search for a Motor Program

Experiments with tethered flies allow the study of simple motor outputs other than yaw torque. *Drosophila* can operate arena rotation (flight direction) also by lift/thrust, leg

¹Present address: Department of Neurobiology and Anatomy, The University of Texas — Houston Medical School, MSB 7.168A, Houston, Texas 77030, USA.

²Corresponding author.

E-MAIL heisenberg@biozentrum.uni-wuerzburg.de; FAX 49-931-8884452.

Article and publication are at www.learnmem.org/cgi/doi/10.1101/lm.37501.

posture, or abdomen position (Fig. 1; Wolf and Heisenberg 1991; Wolf et al. 1992). As the fly initially has no clue as to which behavior the experimenter chooses for control of the arena movements, the animal has no choice but to activate its repertoire of motor outputs and to compare this sequence of activations to the dynamics of arena rotation until it finds a correlation. This is, indeed, what happens. No matter to which side the fly has to shift the abdomen or legs in order to cause clockwise (or, respectively, counterclockwise) rotation, it manages to stabilize the panorama. The fly is “trying out” its behavioral repertoire. This is called operant behavior.

These findings reveal an important property of behavioral flexibility. The visuo-motor link cannot be permanently installed but must form temporarily for the particular task. In operant behavior, the behavior is not a response. The animal generates it to manipulate the stimulus. Only as a result of trial and error is the visuo-motor link firmly closed (stimulus → response). The behavioral activity in operant behavior has been called “initiating” activity (Heisenberg 1983).

This internal search for the effective motor program is not entirely random. The fly is evolutionarily predisposed to interpret large-field visual motion as self motion and tries to counteract it by syndirectional yaw torque (the well-known

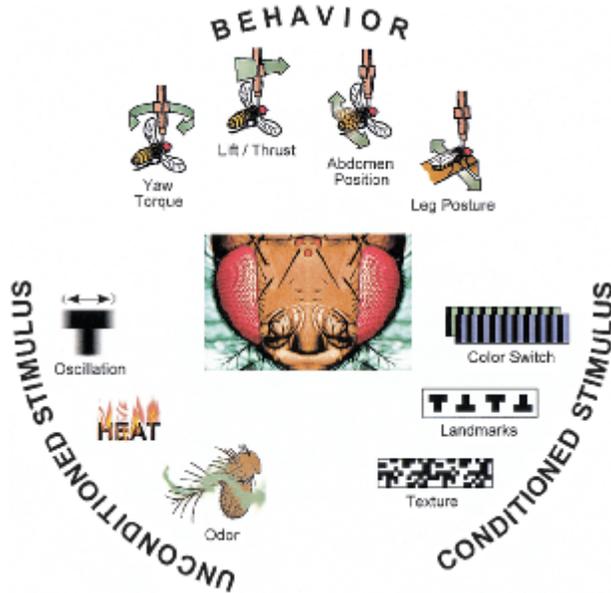


Figure 1 Tethered flies display many kinds of behavioral plasticity. They can use any of a variety of behavioral outputs (top) to control important sensory stimuli (bottom left; unconditioned stimuli [US]). Four behavioral variables (yaw torque, lift/thrust, abdomen position, leg posture) are mentioned in the text; only yaw torque is discussed in detail. A US can be associated with all kinds of (originally neutral) stimuli (bottom right; conditioned stimuli [CSs]) if these stimuli have a predictive value for the US. Often, learning experiments comprise all three components: a behavior, a CS, and a US (three-term contingency).

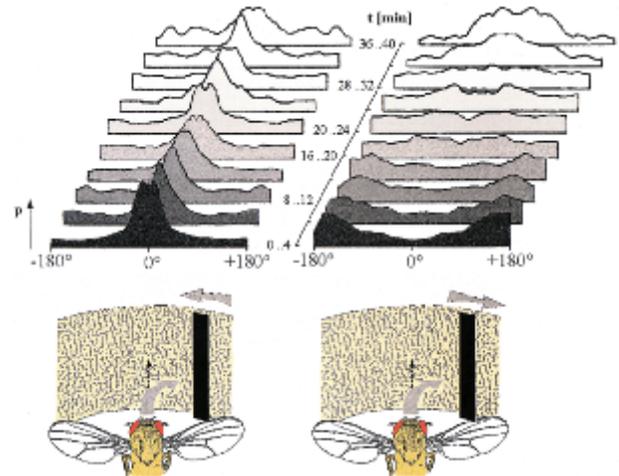


Figure 2 The fly is flying in the flight simulator surrounded by two concentric cylinders. The outer one carries a texture, the inner one a single vertical stripe. (Left side) In normal operation, the panorama with the texture and the stripe turns left if the tethered fly tries to turn right, and vice versa. The upper section shows histograms of flight direction with respect to the stripe of 10 consecutive 4-min periods. Flies have a tendency to head toward the stripe. (Right side) If the coupling is inverted for the stripe, the fly first shows antifixation but learns in the course of 40 min to shift the stripe to the front and to stabilize it there.

optomotor yaw response). As a consequence, if we switch the coupling of arena rotation from lift/thrust to yaw torque, it takes the fly <200 msec to resume control. If we switch from yaw torque to lift/thrust, the fly needs several seconds. Apparently, on encountering arena rotation, the fly is likely to try syndirectional yaw torque before trying other behaviors (Heisenberg and Wolf 1984; Wolf and Heisenberg 1991).

Two experiments, both involving yaw torque, are particularly informative in illustrating this basic organization of the visuo-motor interface. In the first experiment, the fly is surrounded by two concentric cylinders. The outer one carries a texture just coarse enough to be resolved by the fly’s eye. The inner one is transparent and carries one vertical black stripe (Fig. 2). Tethered flies have the tendency to head toward the stripe. An investigation of this arrangement shows that stripe and texture mediate different aspects of the behavior. The texture allows the animal to fly straight and the stripe serves as reference for choosing a particular orientation, as will be further discussed below (Heisenberg and Wolf 1984, 1993).

In the critical part of the experiment, the polarity of the coupling for the stripe is inverted. The turning commands (torque spikes), normally shifting the stripe to the front, now move it backward. By trying to correct this unexpected motion, the fly causes a catastrophic acceleration of the movement toward the back. Individual flies try to cope with this situation in different ways and need different amounts of time to find a solution. Eventually, however, in

the course of 45 min, every fly manages to stabilize the stripe in frontal positions. Not only does the fly learn to use torque spikes of opposite polarity for shifting the stripe, it also has to reorganize the effects of the efference copies of these torque spikes (Heisenberg and Wolf 1984).

The second experiment is performed with a genetic variant lacking most of the optic lobes. These flies still show phototaxis and, at the torque meter, respond to the sudden shift of a landmark by a burst of arbitrary polarity but generate no syndirectional yaw torque in response to horizontal pattern motion. Nevertheless, they can stabilize the panorama in the flight simulator. Interestingly, this performance is independent of the coupling polarity. The flies stabilize the panorama even if yaw torque and pattern motion are positively coupled (i.e., yaw torque to the right causes pattern motion to the right and vice versa). This excludes the involvement of directional optomotor responses in the stabilization (Heisenberg and Wolf 1984). The mutant flies must perceive the amount (strength) of motion (Wolf and Heisenberg 1986) and stabilize the panorama by trying out behaviors.

Assuming that a stable flight orientation is a desired state and that these partially blind animals still take the (nondirectional) visual motion they perceive as self motion, we can explain their behavior by the following rule: Continue a behavior that reduces, and abandon a behavior that increases, the deviation from the desired state. This rule explicates operant behavior (see above). With it, the appropriate behavior would be automatically selected if it were initiated among other behaviors. Note that operant behavior turns into operant conditioning if a lasting modification of the chosen motor program ensues. The above rule in itself is not suited to produce motor learning.

For the discussion below, it may be useful to realize that in operant behavior and classical conditioning, the unconditioned stimulus (US) serves entirely different functions. In both forms of learning, the aversive stimulus (US, heat) is the deviation from a desired state. In classical conditioning the animal is predisposed to respond to it by a fixed behavior, whereas in operant behavior, as studied here, the animal activates its behavioral repertoire and makes the duration of the respective behaviors dependent on the changes in the US.

The two experiments above demonstrate that visual orientation in the flight simulator is an operant behavior. This applies at least to the nondirectional component of motion and to the choice of flight direction by torque spikes. Whether the directional component, the optomotor response, also operates according to these principles or whether it works as a more automatic feedback loop is still a matter of debate (Egelhaaf and Borst 1993).

Landmarks

With the flight simulator, more than just visual feedback is

introduced. A singularity in the arena is a reference for orientation, just like the sun or a distant landmark in free flight. It enables the fly to choose an orientation and to perform certain consistency checks. For instance, continued turning to one side eventually should bring the fly back into the orientation from where it started. This operation can be observed in the flight simulator. Flies frequently generate a train of torque spikes that turn the panorama by $\sim 360^\circ$. These loops can start with the stripe at any angular position (Fig. 54 in Heisenberg and Wolf 1984). A further consistency check is turning to one side followed by the same amount of turning to the other side. This requires measuring the angle turned. It has been demonstrated that the fly in the flight simulator measures its angular distance from the reference orientation by turn integration (Wolf and Heisenberg 1997) — in other words, at any time the fly possesses an orientation vector relative to the reference.

Flight Simulator Learning

We now consider flight simulator (fs) learning in the presence of visual landmarks. In the standard experiment, heat is again the reinforcer.

The arena is decorated with a set of four T-shaped patterns of alternating orientation, two upright and two inverted, each in the center of one quadrant (see Fig. 1, landmarks). Before training, the fly shows a moderate fixation of the patterns without a striking preference for one or the other (Fig. 3). In the training, the punishing heat beam is applied whenever the fly chooses a flight direction in, say, the quadrant marked by an upright T. The fly will immediately perform escape maneuvers and eventually find a 'safe' flight direction toward one of the inverted Ts. It may still make short excursions into the 'hot' quadrants but will more and more effectively avoid the heat. If, after a few minutes of training, the heat is permanently switched off, the fly still prefers flight directions toward the previously safe patterns (Fig. 3). Eventually, however, the learned preference will be extinguished, and the distribution of flight orientations will return to that of the pretraining period.

Pattern memory can be detected 48 h after a spaced training schedule comprising 36 min of training. Pharmacological studies reveal at least three memory phases. Anesthesia-sensitive memory lasts up to 20 min after training. An anesthesia-resistant memory phase follows that is also insensitive to protein synthesis inhibitors. Finally, memory 3 h after training is affected if protein synthesis is prevented (Xia et al. 1997).

Reinforcers

In the terminology of learning psychology, during training the fly transfers the avoidance-eliciting property of the US (heat) to the conditioned stimulus (CS; T-shaped patterns). The patterns indicate, or predict, the heat. We have mentioned above that flies can measure their angular distance

from a reference orientation. Avoidance of the heat is much easier if it is correlated with certain flight orientations. Not unexpectedly, increasing the intensity of the heat also increases both the avoidance scores during training and the learning scores afterwards. Also, other deviations from the desired state are reinforced in the flight simulator. For instance, exafferent motion can be used as the US. If the arena is oscillated with a frequency of 5 Hz and an amplitude $A_{(p-p)} = 15^\circ$ whenever the fly is oriented toward one of the patterns, the fly tries to avoid these orientations and continues to do so later when the oscillations are no longer applied (Wittig 1995). Similarly, one can apply a noxious odor (benzaldehyde) as an aversive US to induce a preference of flight direction relative to visual landmarks (Guo and Götz 1997).

Conditioned Stimuli

In a further variation of this experiment, the arena wall can be decorated with a random dot texture instead of isolated patterns. As is the case with the patterns, the texture has a repetition period of 180° , with opposing quadrants carrying the identical distribution of pattern elements (see Fig. 1, texture). With this type of CS, flies learn to avoid the hot quadrants as well as they do in the experiments with isolated landmarks. Interestingly, irrespective of the particular distribution of pattern elements in the texture, *Drosophila* keeps its flight direction during and after training in the center of the cold quadrants, that is, as far away from the hot quadrants as possible (Dill et al. 1995).

Finally, colors can serve as conditioned stimuli as well. In a panorama that carries identical landmarks in all quadrants, flies can learn to associate colors with heat and no heat if the illumination of the whole arena is made to change color whenever the fly's orientation passes a quadrant boundary.

Learning the Flight Simulator

In contrast to yt learning, fs learning does not involve the modification of motor programs. This is important for understanding the particular operant nature of fs learning and will be discussed in the next section. However, considering how artificial the situation in the flight simulator is, it seems plausible that *Drosophila* would adjust its motor performance to the special feedback conditions. Indeed, if one increases the coupling coefficient between yaw torque and angular velocity of the panorama (i.e., electronically equipping the fly with extra-large wings), the fly needs several minutes of practice without heat before it performs well in the fs-learning task (Guo et al. 1996).

The Relationship between Operant and Classical Conditioning

Both yt learning and fs learning involve operant behavior. Evidently, yt learning is a case of operant conditioning. But what is fs learning? The fly learns to associate certain colors

or patterns with heat and “no-heat,” thus meeting the definition of classical (Pavlovian) conditioning. Could the operant behavior in fs learning merely accompany a learning process that, in essence, is classical? This question is answered by a control experiment. Exact sequences of pattern position and heating periods are recorded during operant training and are played back to a naive fly. Hence, the latter fly receives exactly the same visual and thermal input during the entire training period as did the first fly, but it has no control of the rotations of the panorama and temperature. The “yoked” control is a purely classical conditioning experiment.

It turns out to be considerably less effective than the training operantly controlled by the fly. The operant behavior improves learning. “Do it yourself” would be as good advice for flies as it is for humans.

Thus, fs learning is a composite learning task comprising a behavior (BH), a CS, and a US. The paradigm exemplifies the three-term contingency emphasized by Skinner (1938) and others. The present experiment for the first time provides the opportunity to rigorously dissect the three components. How are operant and classical learning related in a composite learning task? Why is “Do it yourself” so effective?

A simple answer would be that in the three-term contingency, both a classical CS \rightarrow US association and an operant BH \rightarrow US association are formed in parallel and that the two components are summed. Because the classical component can be isolated (yoked control), the critical question concerns the operant component. A direct behavioral modification (as in yt learning) cannot occur, because the same torque spikes bring the fly into, as well as out of, the heat.

However, the fly might learn a strategy such as “stop turning when you come out of the heat.” This improved avoidance behavior would then amplify the effect of the classical CS \rightarrow US association in the memory test. The advantage of the tethered fly at the torque meter is that its yaw torque is available in exquisite detail.

Different avoidance strategies should show in the timing, frequency, size, or polarity of torque spikes after passing a quadrant boundary but were not found (Brembs and Heisenberg 2000). In fs learning, the operant behavior appears to facilitate the acquisition of classical associations. In the classical training, the fly may, for instance, be preoccupied (distracted) by searching for a behavior that would control the temperature.

More can be learned about three-term contingencies by a variant of yt learning called switch-mode (sw) learning (Fig. 3, right side). In this paradigm the two yaw torque domains are not only combined with heat or no-heat, as in yt learning, but also with CSs such as colors or T-patterns.

If, for instance, yaw torque is in the left domain, the panorama is illuminated in, say, green; when yaw torque enters the right domain, the illumination is switched to

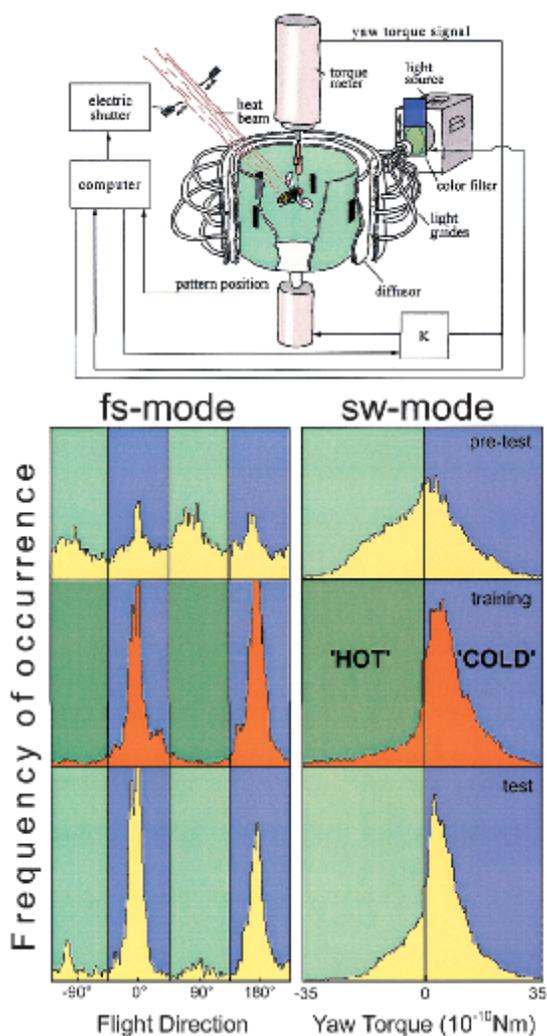


Figure 3 Flight-simulator (fs) and switch-mode (sw) learning. In this experiment, colors (blue and green) are the conditioned stimuli, and in fs mode the quadrants carry identical vertical stripes that give no cue whether or not the quadrant is contiguous with heat. The *upper* diagram depicts the apparatus used in both kinds of experiments (see text for details). Shown in the *lower* panels are representative histograms of flight direction (*left*) and yaw torque (*right*) during pretest, training, and memory test. Shown is the mean of >20 flies in each experiment. Each color is contiguous with heat in half of the experiments. For clarity, heat is depicted as being contiguous with green.

blue. Switch-mode learning causes a larger learning score than yt learning, as the training in fs learning is more effective than in the replay (yoked control). In other words, in both cases, the three-term contingency (CS → US, BH → US) is more efficient than the two-term contingency (BH → US in yt learning; CS → US in the yoked control of fs learning). The interesting difference is that in fs learning, as discussed above, the behavior is not directly modified, while in sw learning the endogenously generated distribution of yaw-torque values is altered, just as in yt learning.

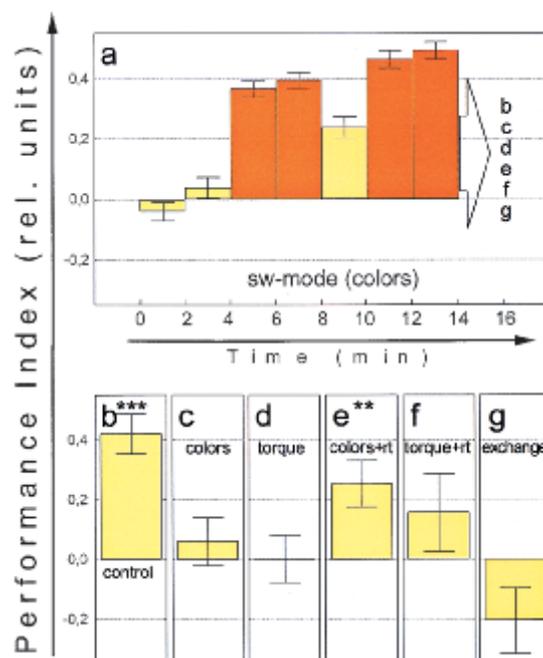


Figure 4 Switch-mode training with colors keyed as in Figure 3. Performance index $PI = (t_a - t_b) / (t_a + t_b)$; with t_a being the time yaw torque is in the cold domain and t_b the time yaw torque is in the hot domain. In *g*, the contiguity between colors and yaw-torque domains is exchanged. A negative PI indicates that the fly prefers the previously cold yaw torque domain and not the cold color. Each bar represents 2 min of measurement. In *e* and *f*, the fly received a short reminder training (rt) in the new situation of the subsequent memory test. This reminder training in itself does not generate a significantly positive PI. In *c* and *e*, the PI refers to flight direction in the flight simulator. Here, t_a and t_b indicate the time flight direction is in the cold and, respectively, hot quadrants. Pretests and memory tests are in yellow; training is in orange. The pretest, first training, first memory test, and second training are averaged (*a*) for all six experimental groups (*b-g*, $n = 250$). (B. Brembs and M. Heisenberg, in prep.)

This allows us to investigate which of the components are learned in the three-term contingency (B. Brembs and M. Heisenberg, in prep.).

If after sw training the CSs are omitted and the endogenously generated yaw torque modulations are tested alone (as in yt learning), the fly, surprisingly, shows no memory at all (Fig. 4d,f) Yet the fly does learn the BH → US association because if, instead, for the test the contiguity between colors and yaw-torque domains is exchanged, no color preference is observed (Fig. 4g). It seems that in sw learning, the behavioral modification is color coded. The fly displays it only in the right color context (B. Brembs and M. Heisenberg, in prep.).

To test the CS → US association alone after sw training, the fly is tested in the flight simulator (i.e., the choice of the fly's flight direction and not its yaw torque determines the color of the arena illumination). The fly still displays the previously acquired color preference in the new paradigm, although to show it, the fly requires a brief reminder train-

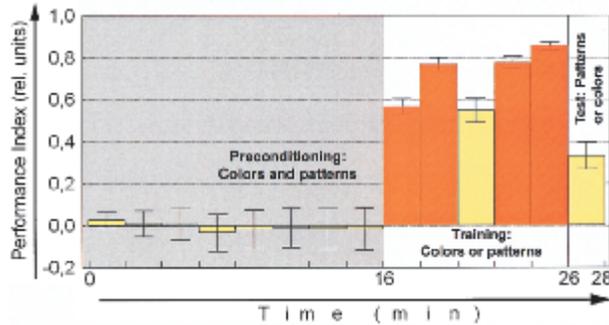


Figure 5 Sensory preconditioning in the flight simulator. The fly is exposed to colors and T-shaped patterns simultaneously without reinforcement. For instance, if the upright T is in front, the arena is illuminated in green; if the inverted T is in front, the illumination is blue. After 16 min, either patterns or colors are removed and the fly undergoes a regular training with heat and the remaining stimuli (either colors or patterns). Finally, memory is tested with the alternative stimuli (i.e., those removed before the training). Bars indicate performance indices (PIs) of 2-min periods, as explained in the legend to Figure 4. $n = 56$. (B. Brembs and M. Heisenberg, in prep.)

ing (Fig. 4e). This remarkable finding requires a multimodal representation of sensory space in the fly brain, where colors and patterns are stored and combined with the good and bad of temperature values, noxious odors, or exafferent motion (B. Brembs and M. Heisenberg, in prep.). As will be discussed below, even colors and patterns are stored and tied to each other if presented together without reinforcement.

To conclude, in both fs and sw learning, only the classical but not the operant association is separately accessible, ruling out that the two act additively. Not summation but interaction between the operant and classical components characterizes the three-term contingency. Apparently, learning by doing works because learning about sensory stimuli is enhanced once the fly can control the stimuli by its own behavior, whereas behavioral modifications tend to be avoided. A long-standing debate about the equivalence of operant and classical processes in learning is finally being confronted with data, at least in *Drosophila*. It would be interesting to test for the generality of these results in other organisms.

Conditioning with Compound Stimuli

Because fs learning works with colors or patterns as CSs, it lends itself to studying the interactions between several CSs and the US. Several of them turned out to be common among most vertebrate species. The question of the extent to which they also apply to invertebrates remains open. Three kinds of interactions were tested: blocking, second-order conditioning, and sensory preconditioning.

In blocking, first one stimulus is trained (CS1 + US) until the subject shows a maximal learning response. Then a new stimulus is added, and the compound is reinforced

(CS1 + CS2 + US). If afterward CS2 is tested alone, the subject usually shows a learning score below that of a control group that has not received the first training. Thus, the first training has blocked learning about CS2. It has been suggested that CS2 is ignored because CS1 already fully predicts the US. Current learning theories incorporate blocking as a critical constituent (Hammer 1997; Schultz and Dickinson 2000).

Attempts to show blocking in fs learning have so far failed, although in many respects the crucial conditions to detect blocking had been met (B. Brembs and M. Heisenberg, in prep.). For instance, at the beginning of the second phase, the CS1 still had a high associative strength in the compound, and no overshadowing occurred between the two CSs during compound training in a control experiment (i.e., colors and patterns accrued about equal associative strength). However, in an ever-changing world, any additional predictor of a reinforcer should be welcome. Perhaps in tethered flight with a single degree of behavioral freedom, the predictive strength of CS1 cannot get high enough. This negative result is in line with the still scarce and debatable reports of blocking in other invertebrates (e.g., Gerber and Smith 1998).

A further interaction possibly taking place during memory acquisition is second-order conditioning (SOC). First the CS1 + US and, in the second phase only, the compound stimuli (CS1 + CS2) without the US are presented. In the final phase, CS2 is tested. Second-order conditioning is said to have occurred if the learning score for CS2 is higher than in a control group that did not receive any training before the compound presentation. In fs learning, weak SOC is observed (B. Brembs and M. Heisenberg, in prep.).

Exchanging phases 1 and 2 in SOC finally yields a sensory preconditioning (SPC) experiment. Here the unreinforced compound presentation (CS1 + CS2) precedes the single CS1 + US training. Sensory preconditioning is said to have occurred if, afterward, CS2 alone shows a significant learning score.

Robust SPC is observed in the flight simulator after 16 min of compound presentation (Fig. 5; B. Brembs and M. Heisenberg, in prep.). SPC is a case of incidental learning where two neutral stimuli are symmetrically associated with each other without any explicit reinforcement. The significance of this finding will become apparent in the next section.

What is the Memory Template?

We now turn to the question of what the fly actually stores in yt and fs learning. During training in the yt-learning paradigm, the fly can avoid the heat by trying out, that is, by the same rule that had worked for the partially blind flies in the flight simulator for stabilizing the panorama: Extend a behavior that reduces, and abandon a behavior that increases, the deviation from the desired state. As stated above, oper-

ant behavior does not imply conditioning. This may arise in two ways. Either the fly modifies the probabilities of its spontaneously generated yaw torque values during the training and keeps this altered distribution beyond the training period or it associates its proprioceptive sensory state of intended turning (e.g., the patterns of excitation from the mechanosensory hairs in the joints, etc.) with the heat during the training phase and, subsequently, avoids this sensory state during the memory test. The former case would be true motor learning. In the latter, motor learning would be based on classical conditioning, and the above rule (operant behavior) would remain the only operant component. So far we cannot distinguish between these possibilities.

Context Generalization

In the remainder of this section, we turn to fs learning. In classical conditioning the fly binds together two types of stimuli, the US and the CS. The US is the signal of an important property of the outside world, for which the animal or its ancestors in phylogeny have discovered (by trying out) a behavior that controls it. All other stimuli are eligible, in principle, to become a CS. A CS indicates or predicts a US. What matters for the selection of a predictor is the degree to which it bears a relation to the US.

Classical conditioning is a pattern-recognition problem. The predictor is a memory template that for retrieval has to be matched with the actual situation. Conceptually, the simplest template is a stored representation of the complete sensory input coinciding with or immediately preceding reinforcement. However, such a predictor has little use in a continuously changing world. If the memory template is very detailed and has to be precisely matched, the system will generate few false positives but will miss many of the events it was supposed to predict. If, however, the matching process is sloppy, it will identify not only many of the relevant events but also many false positives. To avoid this dilemma, the animal needs to make the template as specific as possible. A stimulus that is highly correlated in time with the US is likely to be part of the same state or event in the world. A less tightly (but significantly) correlated stimulus may still have its value as a tentative predictor. Hence, one of the first processing steps to which a brain must subject sensory stimuli for selecting CSs should be a comparison of their temporal structure with that of the US and among each other (see SPC above). Indications of this process have, indeed, been found in fs learning.

In fs learning, the patterns or colors are taken as the CSs. They are the only reliable predictors of heat and no heat, and they control the animals' behavior in the memory test. This does not exclude, however, the possibility that other nondiscriminative stimuli such as the scent and humidity in the room, the stationary light on the ceiling, and the color of the arena illumination might be part of the

memory templates as well. They would not interfere with memory as long as they would not change between training and test. To test this possibility, Liu et al. (1999) have systematically varied features of the arena illumination between training and test. It turns out that with T-patterns as CSs, *Drosophila* can generalize across various contexts (two different color pairs as well as 200-msec 'dark flashes' every 3 sec vs. constant illumination). Thus, the fly does in fact distinguish between the predictive cues (CSs) and the context. It does perform a temporal analysis of the incoming stimuli. However, the fly's ability to generalize across contexts is limited. If the arena illumination during training is sufficiently different from that during the test, memory is disturbed.

Interestingly, context changes that interfere with memory lead to permanent memory loss within 15–30 sec, as can be shown by switching back to the original context (Wiener 2000).

Pattern Recognition

Visual pattern recognition is a prerequisite of fs learning with patterns. A memory template of the pattern has to be generated and stored together with the restrictive or permissive temperature during acquisition and recognized during retrieval. A behaving animal with eye movements of less than the interommatidial angle (Franceschini et al. 1991) provides access to properties of visual pattern recognition that are difficult to explore otherwise. However, it is important to emphasize that the pattern-recognition system so far accessible in *Drosophila* is one of landmark orientation and of tethered flight—not necessarily the entire system.

In *Drosophila*, as in the honeybee, pattern recognition works independent of pattern motion (Srinivasan et al. 1993; Wolf and Heisenberg 1995). Patterns can still be recognized as stabilized retinal images. Early in the investigation it was discovered that in fs learning, *Drosophila* shows no retinal transfer for vertical displacements and between the visual half fields (Dill et al. 1993; Dill and Heisenberg 1995). Moreover, many aspects of pattern shape played no role, whereas the overlap between the memory image and the actual image in relation to the size of the actual image quantitatively described to which extent a modified or displaced pattern could still be recognized. These findings were compatible with a simple mechanism involving a pixel-by-pixel memory template that had to be retinotopically matched to the actual image (pixel matching; Dill et al. 1993). Later it became clear that first, fs learning as an orientation task inevitably requires template matching and, second, that any task that can be solved by pixel matching can also be solved by parameterization if the spatial coordinates on the eye are taken as parameters (for review, see Heisenberg 1995).

Meanwhile, a large number of pattern pairs have been

studied. Some with little overlap can be spontaneously discriminated but are not learned as predictors of heat. In other pairs, conditioned discrimination can be shown to be dependent on parameters unrelated to overlap. On the basis of these results, the pixel-matching hypothesis must be rejected (Wolf and Heisenberg 1998; Ernst and Heisenberg 1999).

To date, four parameters have been identified that the fly can use in conditioned pattern discrimination: vertical position of the center of gravity of the patterns, pattern size, horizontal/vertical extent of the patterns, and vertical separation of pattern elements. The fly already shows a full-sized learning score if the pattern pairs differ in only one of these parameters while all others are kept constant. Apparently, *Drosophila* regards pattern elements that are grouped together as a single figure because its behavior depends on the common center of gravity of these elements. This phenomenon may be related to the experience of humans reporting grouped pattern elements to appear as a single figure (Ernst and Heisenberg 1999).

For one of the parameters (vertical position of the center of gravity), generalization experiments have been performed. Each of the patterns consisted of two separate horizontal bars. Flies were trained to distinguish one pair of patterns and were then tested with a different pair that did not overlap with the first but retained the centers of gravity.

Flies remembered the centers of gravity previously associated with heat and no heat and showed retinal transfer with respect to the pattern elements.

It is important to show as a control that the patterns exchanged in the transfer can be successfully used as a pair in conditioned pattern discrimination. In pattern generalization, as in context generalization, parameters that do not contribute to the discrimination are not part of the memory template. Finally, we note that this retinal transfer is again at odds with the pixel-matching hypothesis (Ernst and Heisenberg 1999).

The fly stores and recalls patterns quickly and without reinforcement. This is demonstrated in an experiment called novelty choice. The fly is exposed during 1 min to four identical patterns in the flight simulator.

Subsequently, in two opposing quadrants the patterns are exchanged with a new one. Aside from its spontaneous pattern preferences, the fly has the tendency to turn toward the new pattern. Interestingly, the exposure phase shows no operant facilitation (Dill and Heisenberg 1995). We have discussed, above, another case of incidental learning: SPC. It is a very different process. Colors and patterns are associated because of their similar spatiotemporal properties with respect to flight orientation. The association takes 16 min (or at least >10 min) of exposure in the flight simulator to form. It is probable that parameter extraction and storage in novelty choice is much faster because it does not require a lengthy temporal analysis.

Selective Visual Attention

A further level of flexibility in yaw torque is selective visual attention. This has been encountered in three experimental situations at the torque meter, and recently, Schuster (1996) has reported potentially related phenomena in walking flies. We will restrict the discussion to selective attention at the torque meter.

If a single vertical black stripe is harmonically oscillated on the side of the fly (mean position $\psi = \pm 45^\circ$; $f = 0.3$ Hz; $A_{p-p} = 30^\circ$), the fly responds to this motion with a characteristic pattern of torque. It directs torque spikes preferentially toward the stripe and, in addition, generates a strong tonic syndirectional torque response at the beginning of the front-to-back phase, while during the back-to-front phase it reduces the torque only gradually to the original level. Now, what happens if two stripes oscillate in antiphase on either side of the fly?

The fly might be expected to show a random pattern of torque modulations because all the positional and rotational components of the stimulus would be perfectly balanced. In contrast, the fly generates for part of the time the typical pattern for the left stripe, at other times that for the right stripe, and again at other times a random pattern. The mean yaw torque is in the left domain when the pattern for the left stripe is generated and vice versa. Apparently, the fly can restrict the processing of visual stimuli to part of the visual field and, while trying to turn to one side, restrict its processing window to the side that is likely to move toward the frontal position (Wolf and Heisenberg 1980).

In the flight simulator it is possible to observe a similar phenomenon with two stripes, one of which is suddenly displaced with respect to the other and subsequently slowly returned to the previous angular distance. With a single stripe, such a displacement leads to a phasic syndirectional response in ~80% of the cases. In the presence of a second stripe, this frequency falls to just half this value, as if the processing window was turned to the displaced stripe for only half of the time. The experiment shows, further, that the processing window often switches sides immediately after the fast displacement. The fly ignores the stripe during the fast displacement but stabilizes it during the slow return (Heisenberg and Wolf 1984).

Above, a further variant of yt and sw learning has been briefly mentioned. In this paradigm the yaw-torque domains (representing left or right turns) are contiguous with heat and no heat, and during the entire training period the fly sees, let us say, an upright T on the left and an inverted T on the right. If during the subsequent test period the patterns remain as before, one obtains a robust learning score similar to that of yt learning. If, however, the two patterns are exchanged between training and test, mean yaw torque gradually shifts to the domain previously contiguous with heat. Again, this experiment implies that during training the fly directs its processing window to the side to which it

turns to get out of the heat. Subsequently, the pattern previously viewed in the processing window signals the safe side, and this information slowly overrides the motor learning.

Anatomy and Physiology

How does the brain accomplish these various flexibilities? To what extent can any of the behavioral functions described above be assigned to brain regions, neuronal circuits, neuromodulators, or biochemical processes? The humble answer is: Hardly at all. The little we know relates mainly to fs learning.

Fs learning requires intact genes *rutabaga* and *dunce* that regulate cAMP levels in intracellular signaling. They influence the plasticity of the neuromuscular junction as well as the branching patterns of neuronal processes during development and are known to affect most processes of behavioral plasticity.

At the level of brain structure, two neuropil regions in the central brain have received some attention in recent years, the so-called mushroom bodies and the central complex. Flies without functional mushroom bodies can be obtained by genetic and pharmacological methods. While it is well established that the mushroom bodies are required for olfactory learning and are, at least in part, the site of the short-term memory trace for odors, they are dispensable for both fs learning and the classical replay learning as well as for yt learning. They are required, however, for context generalization as described above. For pattern generalization, however, they are again irrelevant. Other functions such as middle- and long-term memory, SPC, novelty choice, and selective attention have not been tested in mushroom body-less flies to date.

While the influence of the central complex on fs learning may be more direct than that of the mushroom bodies, the experimental base for this conclusion is more tenuous because it is not yet possible to generate the same well-defined lesions in the central complex by independent methods.

Several structural mutants of the central complex are impaired in fs learning. For instance, the mutant *nobridge*^{KS49} (*nob*^{KS49}) lacking the middle part of the protocerebral bridge is impaired in fs learning. Only prolonged training of 4 × 4 min produces a small but significant learning effect. Heat avoidance is little impaired. *nob*^{KS49} flies fail in classical pattern learning and yt learning (Weidtmann 1995). Also, blocking certain fiber systems of the central complex by gene-transfer techniques disturbs fs learning. These methods of intervention that are currently being designed will dramatically advance the kind of analysis presented in this review.

Conclusions

The flexibility of yaw torque in tethered flight consists of at least 10 components:

- (1) Trying out (operant behavior): The fly has to find the behavioral output that controls a significant sensory variable.
- (2) Motor learning: The fly modifies an endogenous motor program according to its experience during the trying-out phase and retains this modification afterwards.
- (3) Operant facilitation: The CS and the US are more easily associated if they both are controlled by the fly's yaw torque (learning by doing).
- (4) A multimodal sensory representation: The CS → US association is independent of behavior.
- (5) Context generalization: Stimuli are temporarily stored and analyzed according to their temporal structure. For a strongly response-related stimulus (US), stimuli with a similar temporal structure are predictors (CSs); all others are context.
- (6) Sensory preconditioning: Two stimuli can be associated even if neither is related to a response.
- (7) Three-term contingencies: In conditioning experiments involving a CS, a US, and a behavior, the component associations interact; motor learning is suppressed.
- (8) Early visual memory: Visual patterns are temporarily stored in a fast, unreinforced process.
- (9) Pattern recognition: Visual patterns are stored as sets of parameter values. Four such parameters have been identified. Only parameters that contribute to the pattern discrimination task are stored. Pattern elements spatially grouped are evaluated as one figure.
- (10) Selective attention: Flies can restrict their visual processing to parts of the visual field. The processing window can be shifted quickly and is preferentially shifted to the side to which the fly tries to turn.

Behavioral flexibility is a large mixed bag. As a key component of flight orientation, yawing may have access to much of the brain's functional complexity. However, viewing this complexity through a single output variable may inspire scientists to eventually place the many isolated descriptions of paradigms into a coherent behavioral model of brain function.

REFERENCES

- 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.
- Dill, M. and Heisenberg, M. 1995. Visual pattern memory without shape recognition. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 349: 143-152.
- Dill, M., Wolf, R. and Heisenberg, M. 1993. Visual pattern recognition in *Drosophila* involves retinotopic matching. *Nature* 365: 751-753.
- Dill, M., Wolf, R. and Heisenberg, M. 1995. Behavioral analysis of *Drosophila* landmark learning in the flight simulator. *Learn. Mem.* 2: 152-160.
- Egelhaaf, M. and Borst, A. 1993. A look into the cockpit of the fly: Visual orientation, algorithms, and identified neurons. *J. Neurosci.* 13: 4563-4574.
- Ernst, R. and Heisenberg, M. 1999. The memory template in *Drosophila* pattern vision at the flight simulator. *Vision Res.* 39: 3920-3933.

- Franceschini, N., Chagneux, R., Kirschfeld, K. and Mücke, A. 1991. Vergence eye movements in flies. In *Synapse-transmission-modulation: Proceedings of the 19th Göttingen Neurobiology Conference* (eds. N. Elsner and H. Penzlin), p. 275. Georg Thieme, Stuttgart.
- Gerber, B. and Smith, B.H. 1998. Visual modulation of olfactory learning in honeybees. *J. Exp. Biol.* **201**: 2213–2217.
- Götz, K.G. 1964. Optomotorische Untersuchung des visuellen Systems einiger Augenmutanten der Fruchtfliege *Drosophila*. *Kybernetik* **2**: 77–92.
- Guo, A. and Götz, K.G. 1997. Association of visual objects and olfactory cues in *Drosophila*. *Learn. Mem.* **4**: 192–204.
- 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.
- Hammer, M. 1997. The neural basis of associative reward learning in honeybees. *Trends. Neurosci.* **20**: 245–252.
- Heisenberg, M. 1983. Initiale Aktivität und Willkürverhalten bei Tieren. *Naturwissenschaften* **70**: 70–78.
- . 1995. Pattern recognition in insects. *Curr. Opin. Neurobiol.* **5**: 475–481.
- Heisenberg, M. and Wolf, R. 1984. *Vision in Drosophila: Genetics of microbehavior*. Springer, Berlin.
- . 1993. The sensory-motor link in motion-dependent flight control of flies. *Rev. Oculomot. Res.* **5**: 265–283.
- Liu, L., Wolf, R., Ernst, R. and Heisenberg, M. 1999. Context generalization in *Drosophila* visual learning requires the mushroom bodies. *Nature* **400**: 753–756.
- Schultz, W. and Dickinson, A. 2000. Neuronal coding of prediction errors. *Annu. Rev. Neurosci.* **23**: 473–500.
- Schuster, S. 1996. "Repräsentation visueller Objekte beim Suchlauf der Fliege *Drosophila*." Doctoral thesis, University of Tübingen.
- Skinner, B.F. 1938. *The behavior of organisms*. Appleton, New York.
- Srinivasan, M.V., Zhang, S.W., and Rolfe, B. 1993. Is pattern vision in insects mediated by cortical processing. *Nature* **362**: 539–540.
- Weidmann, N. 1995. "Visuelle Flugsteuerung und Verhaltensplastizität bei Zentralkomplex-Mutanten von *Drosophila melanogaster*." Diploma thesis, Julius-Maximilians Universität Würzburg.
- Wiener, J. 2000. "Kontext-Generalisierung in *Drosophila melanogaster*." Diploma thesis, Julius-Maximilians Universität Würzburg.
- Wittig, T. 1995. "Visuelles und motorisches Lernen von *Drosophila*-Fliegen ohne Pilzkörper." Diploma thesis, Julius-Maximilians Universität Würzburg.
- Wolf, R. and Heisenberg, M. 1980. Fine structure of yaw torque in visual flight orientation of *Drosophila melanogaster*. 2. A temporally and spatially variable weighting function for the visual field and visual attention. *J. Comp. Physiol. A* **140**: 69–80.
- . 1986. Visual orientation in motion-blind flies is an operant behavior. *Nature* **323**: 154–156.
- . 1991. Basic organization of operant behavior as revealed in *Drosophila* flight orientation. *J. Comp. Physiol. A* **169**: 699–705.
- . 1995. Learning of *Drosophila* in the flight simulator: Classically conditioned visual pattern discrimination. In *Nervous systems and behaviour: Proceedings of the 4th International Congress of Neuroethology* (eds. M. Burrows, et al.), p. 184. Georg Thieme, Stuttgart.
- . 1997. Visual space from visual motion: Turn integration in tethered flying *Drosophila*. *Learn. Mem.* **4**: 318–327.
- . 1998. Various mechanisms for visual pattern memory in *Drosophila melanogaster*. In *Fifth International Congress of Neuroethology*, (eds. B. Kristan and H. Karten) San Diego.
- Wolf, R., Voss, A., Hein, S., and Heisenberg, M. 1992. Can a fly ride a bicycle? Discussion on natural and artificial low-level seeing systems. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **337**: 261–269.
- Xia, S.Z., Liu, L., Feng, C.H., and Guo, A. 1997. Memory consolidation in *Drosophila* operant visual learning. *Learn. Mem.* **4**: 205–218.

Received October 27, 2000; accepted in revised form December 20, 2000.