

Outcome learning, outcome expectations, and intentionality in *Drosophila*

Martin Heisenberg

Rudolf Virchow Research Center, University of Wuerzburg, Wuerzburg D-97080, Germany

An animal generates behavioral actions because of the effects of these actions in the future. Occasionally, the animal may generate an action in response to a certain event or situation. If the outcome of the action is adaptive, the animal may keep this stimulus–response link in its behavioral repertoire, in case the event or situation occurs again. If a responsive action is innate but the outcome happens to be less adaptive than it had been before, the link may be loosened. This adjustment of outcome expectations involves a particular kind of learning, which will be called “outcome learning.” The present article discusses several examples of outcome learning in *Drosophila*. Learning and memory are intensely studied in flies, but the focus is on classical conditioning. Outcome learning, a particular form of operant learning, is of special significance, because it modulates outcome expectations that are operational components of action selection and intentionality.

As an agent and a highly autonomous being, an animal generates its behavior to interact with the outside world. The effects of a behavior can be good or bad, adaptive or nonadaptive. A mistake can be fatal. An animal activates a behavior because of its effects. At the time of action selection, these lie in the future. The more behavioral options an animal has, the more difficult it gets for the animal to select the right one. For this demanding task an animal needs its brain. This basic property of behavioral control is little understood.

Occasionally an animal activates a behavior because of an event that has just arrived in the outside world. This is called a reflex or response. In behavioral science, especially with insects, such stimulus–response situations have extensively been studied since they are easily amenable to experimentation. Typically, a stimulus (event) would be presented to the experimental animal and the behavior of the animal in the time following the event would be recorded. If in many or all animals of this kind the same behavior would occur after the same or a similar event, the behavior would be called a response.

There must be a reason why the animal has linked a behavior to a certain event or situation. The behavior must have had a positive outcome in similar situations in the past, the past in the life of the animal or its ancestors. If a response (R) to a certain stimulus (S) would always be adaptive it could be fixed. The S–R link would not have to be modifiable in its strength. As the future is open and the effects of an action in most cases are not taken for granted, the animal is well advised, however, to monitor the outcome of its behavior and assess its adaptive value. This is where outcome learning and outcome expectations come in.

Outcome learning is a variant of operant learning. Operant learning is said to be learning by punishment and reward. The animal responds to an event with a behavior depending on whether its outcome will be punishing or rewarding. As during the activation of the behavior the outcome lies in the future, the animal must rely on its expectation of the putative outcome.

One might be inclined to restrict the term “expectation” to conscious expectations. However, often we become conscious of expectations in retrospect and we know that they were real expectations but we were unconscious of them while they happened. In the following sections, I will discuss the role of outcome learning

and outcome expectations in *Drosophila* reviewing several examples that deal with the plasticity of the stimulus–response link. Other forms of operant learning such as some kinds of sensory adaptation or motor learning may do without outcome expectations and will not be considered here.

For a better understanding of outcome learning we consider its role in action selection more broadly. An action, a behavioral module is a motor program, a score of precisely timed muscle activations driven by a central pattern generator (CPG). Each CPG must have its operating circuit in the brain, which organizes the activation and suppression of this particular action. To do this, it integrates many different kinds of influences, such as the state of the effectors that would have to carry out the action, the external conditions, the nutritional state of the organism, any special sensory signals that might guide the action or might necessitate its immediate activation, motivational drives changing the urgency of its activation, etc. Most important, the operating circuit must assess the consequences of the action once executed. If these have been positive in the past and if a similar situation arises again, this behavior may have a higher probability to be activated. This anticipatory function of the operating circuit in the brain has been called “outcome expectation.”

Let us look at the formation of S–R links. The animal faces a significant event (stimulus S) to which it has not yet assigned an adaptive behavior (response R). In other words, none of the operating circuits have an outcome expectation for this situation in store that would promise to be sufficiently adaptive. At the same moment, however, the animal happens to generate a behavior in an unrelated context or spontaneously. If the behavior that is coincident with S is followed by an adaptive outcome, the operating circuit will generate a memory of the positive outcome of R as an answer to S. A new stimulus–response (S–R) link is born. In case of a similar event (S) the memory of the positive outcome in this operating circuit will favor the activation of the respective behavior (R).

© 2015 Heisenberg This article is distributed exclusively by Cold Spring Harbor Laboratory Press for the first 12 months after the full-issue publication date (see <http://learnmem.cshlp.org/site/misc/terms.xhtml>). After 12 months, it is available under a Creative Commons License (Attribution-NonCommercial 4.0 International), as described at <http://creativecommons.org/licenses/by-nc/4.0/>.

Corresponding author: heisenberg@biozentrum.uni-wuerzburg.de
Article is online at <http://www.learnmem.org/cgi/doi/10.1101/lm.037481.114>.

While no “hard-wired” S–R links are known in *Drosophila*, it cannot be excluded that they exist. They would have no need to be modulated ever and would not require an outcome expectation. In most cases, however, the strength of an S–R relationship in the process of selection of the respective behavior can be adjusted according to experience. This adaptation involves outcome learning and outcome expectations. The memory of the outcome of the behavior is stored and used to update the outcome expectation for the next event.

Case studies

Conditioning of leg posture

Flies can be conditioned to activate and change a particular posture of their legs (Booker and Quinn 1981; Mariath 1985). In one such experiment a tethered fly is placed on a platform that it can shift sideways with its legs. The position of the platform is monitored. The position signal can be used to switch a dangerous heat source on or off. In this way, the fly can operate the switch by displacing the platform with its legs to one or the other side. The fly quickly finds out that it is in control of the heat and keeps the platform for most of the time out of the range where heat is on. At the same time, it still tries to escape the tethering and occasionally tests whether the heat switch is still in operation.

What happens from the fly’s perspective in this behavior? The fly suddenly encounters dangerous heat. With a certain change in leg posture the heat goes off. The operating circuit for leg posture stores the event because of the coincidence between this particular action (change of posture) and the switch-off. When heat occurs the next time, the fly activates this change in leg posture again. Each time the action is again successful, the probabilistic weight of the link between heat and change in leg posture is strengthened. If the success is lacking, the strength of the link is reduced. In this way, a new S–R link tuned to the presumable success of the reaction is formed in this restrained situation on the platform. The fly has established by outcome learning an outcome expectation for a particular action: change of leg posture in response to heat.

Yaw-torque learning

A similar behavior can be observed in flight (Wolf and Heisenberg 1991). The tethered fly is suspended at a torque meter. Yaw torque is recorded in stationary flight. While tethered the fly continuously modulates its yaw torque although these manoeuvres have no effect on the fly’s orientation in space because even the head is glued to the thorax. We measure the fly’s range of spontaneous yaw torque modulations and set its zero point to the middle of that range. The fly again is made to control the heat. Heat is switched on if the fly generates yaw torque, say, to the right (clockwise, cw), heat goes off, if yaw torque is to the left (counterclockwise, ccw). Note that this relation between temperature and yaw torque is entirely artificial and would hardly ever occur in free flight. Still, the fly quickly learns that its yaw torque is in control of the heat and it therefore generates more yaw torque to the left than to the right. Even after the heat is switched off for good (memory test), the fly still prefers for a while to generate yaw torque to the left. It has labeled left turns “safe” and right turns “dangerous” in the context of heat. It seems to expect that yaw torque to the left may prevent and yaw torque to the right may generate heat, even if acutely the latter is not the case. The expectation overrides the new evidence.

Flight simulator learning

At the torque meter the fly can be tested in a so-called flight simulator, where yaw torque drives the angular motion of a panorama

surrounding the fly. Yaw torque to the left causes clockwise rotation of the panorama, yaw torque to the right counterclockwise rotation (“negative” visual feedback of turning). The panorama carries four visual patterns in the centers of its four quadrants (Q1–4), say an upright and an inverted T in alternating sequence. The fly is heated by a laser beam if it heads toward one of the patterns (e.g., the upright T). It quickly learns to avoid the orientations toward quadrants with the upright T, even after the heat is switched off for good (Wolf and Heisenberg 1991).

To show the need for outcome expectations in flight simulator learning one can replace the visual patterns by two colors, e.g., blue and green. The fly is suspended in the center of a striped drum (pattern wavelength $\lambda = 18^\circ$). As before, the drum is conceptually divided into the four quadrants (Q1–4). Depending on its orientation in the drum the fly can switch the background color of the panorama together with heat/no-heat. For instance, if the fly heads toward Q1 or Q3, heat is on and the arena light in all four quadrants is blue. If it heads toward Q2 or Q4, the arena light is green and temperature normal. There are no figural cues to distinguish the quadrants. Nevertheless, while in the “cold” quadrant the trained fly reduces the frequency of turns toward the quadrant boundary depending upon how close its orientation is to that boundary (Wolf and Heisenberg 1997). It must have different outcome expectations with respect to heat for its right and left turns, depending upon its orientation in the quadrant.

Reafferent behavior reveals outcome expectations

Manipulating the motion feedback in the flight simulator we can make the fly show its outcome expectation. The fly is again surrounded by a regularly striped drum and drives the angular motion of the drum with its yaw torque in a negative feedback loop. For flying straight the fly generates long sequences of alternating right and left turns called mini-saccades, each one lasting ~ 100 msec. If we block or accelerate the motion of the drum for 50 msec, not much happens. If, however, we invert the feedback to make it positive for 50 msec, the fly counteracts this disturbance with a saccade of the same polarity as the one that did not give the expected effect (Wolf and Heisenberg 1990). The same is observed during normal saccades. A short (200 msec) inversion of the feedback elicits an immediate correction saccade, whereas other disturbances in this time window have hardly any effect (Heisenberg and Wolf 1993).

Already in the mid-last century, von Holst and Mittelstaedt (1950) described this effect in freely walking flies (*Eristalis*). They ruled out the simple explanation that the fly might just inhibit the visual input during saccades. Instead, they proposed that the brain generates an efference copy of its turning command, calculates from it an outcome expectation and compares it to the reafferent visual motion caused by the turn. If the direction of the reafferent motion does not match the expectation, the fly generates a vigorous response. Analysis in the flight simulator shows that what matters for the comparison is only the direction, not the amount of the reafferent motion. This may be due to the highly artificial feedback situation at the torque meter and the fly’s limited experience with it. Alternatively, it might indicate the fly’s limited ability to calculate outcome expectations.

Learning to expect positive visual feedback

In the flight simulator flies can learn to navigate with a landmark that has its angular motion inversely coupled to yaw torque (positive instead of normal negative feedback). The stripe moves in front of a texture covering the entire inside of the drum. The texture is in normal negative feedback with yaw torque. When the feedback for the single stripe in front of the texture is inverted, the fly generates catastrophic responses to the unexpected

direction of stripe motion but soon manages to suppress them and eventually regains control of it. At an early stage this ability is still fragile. The fly responds to a sudden disturbance of orientation with the wrong correction manoeuvre. Gradually the fly even manages these situations. If after a long period of training the coupling between landmark motion and yaw torque is switched back to normal, for a short while the fly again shows catastrophic responses, as had been observed at the beginning of the experiment (Heisenberg and Wolf 1984). This is the most compelling proof of outcome expectations. During the phase of inverted coupling the fly first suppresses the old outcome expectation, then discovers the appropriate behavior to cope with the new situation and finally generates a new outcome expectation for the turns in the artificial feedback situation.

Habituation may represent outcome learning

Several examples of adaptive behavior categorized as habituation can be better understood as the waning of an outcome expectation. Take the habituation of the landing response (Fischbach 1981; Waldvogel and Fischbach 1991). In tethered flight the fly lifts its front legs in response to an approaching visual object. With repeated stimulation the response probability declines. As no collision and therefore no landing occur, the fly more and more often suppresses the response to the visual stimulus.

Where in the brain is the response probability adjusted? Most likely, it is not at the visual input stage. The visual stimulus used in this study is front to back motion of a black, vertical stripe in one of the visual half-fields (Waldvogel and Fischbach 1991). This stimulus also elicits a phasic yaw-torque response (Sareen et al. 2011) that does not decline in frequency at the same rate as leg lifting. It cannot be motor fatigue either, since a high response probability is restored instantaneously with a switch of the stimulus to the contralateral side. Most likely, the adjustment occurs in the corresponding operating circuit in the brain.

To be useful, any S–R unit needs to fulfill two criteria: it has to be general enough to still work properly if the original event reoccurs in a slightly different context and it must be specific enough to exclude responses to stimuli which are similar but do not represent the event to which the response applies. In the present experiment the operating circuit for leg lifting seems to have a fairly general outcome expectation dealing with approaching objects. The fly learns from the cases in which it does not respond, that this stimulus does not represent the kind of event for which the stimulus–response link had been set up and hence the response does not have a positive effect in this situation. The fly further reduces the response frequency to this particular stimulus. Given that the response frequency for stimuli in the other visual half field stays high, the fly, in effect, increases the specificity of the stimulus in the S–R unit for “landing in response to an approaching object.”

Outcome learning in the heat box

Outcome expectations can also be observed in freely walking flies. Consider heat escape behavior (Yang et al. 2013). A fly is walking in a dark, narrow alley (heat box). Suddenly the heat goes up. The fly turns around and walks in the opposite direction. The heat might be a local event. If so, turning around would be a valid strategy. If the fly learns that turning around and walking into the opposite direction does not terminate the heat, it will suppress the behavior at subsequent similar events. If the fly is hit by the heat pulse during rest, it immediately starts walking (escape) but as in the previous situation it abandons this behavior as soon as it finds out that the behavior does not switch off the heat. The responses are innate and probabilistic. The fly adjusts its corresponding outcome expectations to the outcome in previous events.

Aversive phototaxis suppression

A further example of an outcome expectation is observed in aversive phototaxis suppression (APS) (Seugnet et al. 2009), a behavior that had been part of the original odor avoidance learning that initiated learning/memory research in *Drosophila* (Quinn et al. 1974). To measure APS the fly is placed into a T-maze where it has a choice between two alleys each one leading into a larger compartment, one brightly illuminated and having its walls covered with an aversive chemical (quinine solution on filter paper), the other dark and dry. A naive fly chooses the alley leading to the light with a probability $P > 0.5$. It is positively phototactic. Due to its encounter with quinine at the end of the bright alley, however, its expectation in the next round at the choice point for the outcome at the end of the bright alley is less positive.

What makes this learning special is, that both potential outcomes, quinine and escape, are delayed by at least the passage through the alley. The decision which alley to enter, the lighted or the dark one, is influenced by the expectation of the probable consequences waiting at the other end of the alley, at some time in the future.

APS is a good example showing how important outcome expectations are in behavioral sequences. The fly runs toward the light because it expects to be able to escape where the light is. How do we know? Because, if it is unable to fly for any possible reason such as amputation or gluing of the wings, wingless or flightless mutants, etc., the fly does not run toward the light (McEwen 1918).

Relations between classical and outcome learning

The yaw-torque learning paradigm described above can be modified to reveal a process of classical (Pavlovian) learning. For this we combine the switch between heat and no heat again with the color of the arena light (Brembs and Heisenberg 2000; Brembs 2009). Blue is on when heat is on and yaw torque is to the right; green is on, when heat is off and yaw torque is to the left. This variant has been called switch-mode learning. The additional sensory cue improves the performance in the memory test. Interestingly, in this case the fly forms a color-specific heat–yaw-torque response. If for the memory test the colors are switched off and the fly is tested in a white arena, there is no yaw-torque bias to the left. In switch-mode learning the fly has a positive outcome expectation for yaw-torque to the left, but follows this expectation only if right and left turns are coupled with blue and green. The outcome expectation is not followed in white light. The fly attaches a negative valence to the blue color (S_2) that is contingent with heat (S_1), in other words, it forms an association between blue and heat ($S_2 - S_1$). We can demonstrate this by testing the fly after switch-mode training in the flight simulator as described above. Now the fly can switch between the two colors by choosing its orientation in the panorama. In this new test situation the fly avoids the color that had been combined with heat in switch-mode learning.

In yaw-torque and switch-mode learning during training the modification of the behavior seems to be the same. The fly can switch off the heat or the combination of color and heat by yaw torque to the left. However, in yaw-torque learning in the test phase after conditioning the memory effect is due to a positive outcome expectation for “yaw-torque-to-the-left” behavior, whereas in switch-mode learning the outcome expectation is specific for colored light, due to the formation of the $S_2 - S_1$ association, i.e., the negative valence attached to the blue color previously combined with heat. The color is a decisive qualifier for the outcome expectation that yaw-torque-to-the-left behavior can protect against heat. The fly does not expect the same in white light. The memory test without heat is an extinction training. It will be interesting to find out whether in switch-mode the

memory test affects just the outcome expectation for the behavior in this particular test or the color–heat association itself and consequently also other possible memory tests.

As discussed above, in the flight simulator a fly can be conditioned by heat to keep its orientation away from a certain visual pattern (e.g., upright or inverted T). If during the entire training the background illumination is green, the fly afterward in the test does not show the conditioned pattern preference in blue light. In other words, the outcome expectation is specific for the color the fly had seen during training. Although in this case the color serves only as context and cannot be used by the fly to distinguish between “hot” and “cold” quadrants, the fly generates an outcome expectation that is specific for this color (Liu et al. 1999; Brembs and Wiener 2006). Interestingly, for many color pairs the outcome expectation has no color specificity, the fly generalizes for the color context. Even for such color pairs the fly can generate color-specific outcome expectations with the appropriate reinforcement. In an elegant experiment Brembs and Wiener (2006) showed that flies can learn to avoid pattern A in one color and pattern B in the other.

As in switch-mode learning, also in APS the fly might form a classical S_2-S_1 association, in this case between the light and the bitter taste of quinine, attaching a negative valence to the light. If so, the outcome expectation would be modified because of this association. Such a hypothesis could be tested by training flies in APS and testing their phototaxis in other contexts.

Discussion

Outcome learning exploits the repetitiveness of the world. If the animal has activated a behavior that is adaptive for a particular event and a similar event occurs again, the same behavior should again be adaptive. The various examples above show this causal structure. “Learning by punishment and reward” is often understood as the activation of a behavior because of its outcome. Taken as a statement about a single episode, it would reverse cause and effect. In reality the fly uses its memory of the outcome of the behavior in past events to generate an outcome expectation, i.e., to adapt the action selection process for the release of the right behavior in a similar event to come.

In *Drosophila*, all S–R relationships studied so far have turned out to be modifiable. Therefore, they all must involve outcome expectations. Being based on experience from previous such events they carry estimates of the reliability of the expectation and the adaptive value of the outcome. This is the minimum of what the fly needs for action selection. How detailed and elaborate an outcome expectation is, needs to be investigated in each single case. Once a decision is taken and the actual consequences of this motor act unfold, the fly measures and stores them to update the existing outcome expectation accordingly, to be better prepared for the next incident of a similar kind.

It may be worth pointing out how demanding the use of outcome expectations for action selection may be. The operating center first has to recognize the outcome of its behavior and distinguish it from other processes in the outside world. Second, it must store a working memory of this outcome. Third, it must use this memory to update the outcome expectation derived from the previous events. Updating is needed because events never repeat exactly. Also, the state of the organism or the effectiveness of the executing organs may have changed in the meantime. Moreover, it is often not obvious which parts of the sensory signals represent the significant aspects of the event for the animal. As stated above, the operating circuit for an action must represent the probability of the outcome to occur and what its likely adaptive value might be. It should also determine

how quickly with repeating events a full expectation should build up, how persistent it should be without further confirmation and how quickly it should decay with negative evidence.

Let us at this point recapitulate the differences between the memory traces in classical and outcome learning. Outcome learning modulates the strength of the stimulus–response link according to the outcome in previous such events. The classical switch-mode learning modulates the stimulus–response link via a change in the significance of the stimulus. In both forms of learning the animal adjusts its outcome expectations for the action selection process in the memory test under consideration. In outcome learning the relevant memory traces have been acquired in the past and have served to establish the present outcome expectation for action selection in the present memory test. The memory of the present test will again modify the outcome expectation for future tests. In classical learning the memory trace represents a contingency in the outside world, say, between blue light and heat (S_2-S_1 link). This “knowledge,” however, can only be assessed by a behavioral experiment. The S_2-S_1 link has to be relevant for and modify the outcome expectations in this test behavior and other behaviors regarding heat. Characteristically, after switch-mode learning the fly needs a short reminder training in the flight simulator to increase the probability of its heat avoidance or escape behaviors in the memory retrieval test with blue light (Brembs and Heisenberg 2000).

Outcome expectations are real. They are functional components of action selection and are based on experience. They are represented by neural circuits and membrane properties of neurons in the brain. At the circuit level very little is known about outcome learning and outcome expectations in flies. At the molecular level several genes have been discovered which are required in one or several learning paradigms but not in others. More often than not, this distinction separates operant and classical learning. For instance, the adenylate cyclase gene *rutabaga* is required for classical odor and color learning but not for operant yaw-torque learning (Dudai et al. 1983; Brembs and Plendl 2008) or novelty choice learning (Solanki et al. 2015). It is also dispensable in classical odor intensity learning (Masek and Heisenberg 2008). The genes *pkc* and *dFoxP* are required in yaw-torque learning but dispensable in the classical component of switch-mode learning (Brembs and Plendl 2008; Mendoza et al. 2014). The gene “ignorant” is involved in classical and operant learning in different ways (Putz et al. 2004). These are just a few examples. What they tell us is, that “classical” (world) and “operant” (self) learning are mixed bags, each housing a variety of different behaviors. So far, learning/memory paradigms have focused on the memory process. What we need for each of the paradigms is a better understanding of what the memory does for the behavior and what the behavior does for the animal.

Animals choose between behavioral actions by comparing among other factors the expected future effects of these actions. To adjust their expectations the animals need outcome learning. An action chosen because of its outcome is intentional behavior. Outcome learning and outcome expectations are operational components of intentionality.

Acknowledgments

I thank the vice-president of the Julius-Maximilians-University, M. Lohse for providing laboratory space and my colleagues at the Rudolf Virchow Center for their hospitality.

References

- Booker R, Quinn WG. 1981. Conditioning of leg position in normal and mutant *Drosophila*. *Proc Natl Acad Sci* **78**: 3940–3944.

- Brembs B. 2009. Mushroom bodies regulate habit formation in *Drosophila*. *Curr Biol* **19**: 1351–1355.
- Brembs B, Heisenberg M. 2000. The operant and the classical in conditioned orientation of *Drosophila melanogaster* at the flight simulator. *Learn Mem* **7**: 104–115.
- Brembs B, Plendl W. 2008. Double dissociation of PKC and AC manipulations on operant and classical learning in *Drosophila*. *Curr Biol* **18**: 1168–1171.
- Brembs B, Wiener J. 2006. Context and occasion setting in *Drosophila* visual learning. *Learn Mem* **13**: 618–628.
- Dudai Y, Uzzan A, Zvi S. 1983. Abnormal activity of adenylate cyclase in the *Drosophila* memory mutant *rutabaga*. *Neurosci Lett* **42**: 207–212.
- Fischbach KF. 1981. Habituation and sensitization of the landing response of *Drosophila melanogaster*. *Naturwissenschaften* **68**: 332.
- Heisenberg M, Wolf R. 1984. *Vision in Drosophila*. Springer, Berlin.
- Heisenberg M, Wolf R. 1993. The sensory-motor link in motion-dependent flight control of flies. In *Visual motion and its role in the stabilization of gaze*. (ed. Miles FA, Wallman J), pp. 265–282. Elsevier, Amsterdam.
- Liu L, Wolf R, Ernst R, Heisenberg M. 1999. Context generalization in *Drosophila* visual learning requires the mushroom bodies. *Nature* **400**: 753–756.
- Mariath HA. 1985. Operant conditioning in *Drosophila melanogaster* wild-type and learning mutants with defects in the cyclic AMP metabolism. *J Insect Physiol A* **31**: 779–787.
- Masek P, Heisenberg M. 2008. Distinct memories of odor intensity and quality in *Drosophila*. *Proc Natl Acad Sci* **105**: 15985–90.
- McEwen R. 1918. The reactions to light and to gravity in *Drosophila* and its mutants. *J Exp Zool* **25**: 49–106.
- Mendoza E, Colomb J, Rybak J, Pflüger HJ, Zars T, Scharff C, Brembs B. 2014. *Drosophila* FoxP mutants are deficient in operant self-learning. *PLoS One* **9**: e100648.
- Putz G, Bertolucci F, Raabe T, Zars T, Heisenberg M. 2004. The S6KII (rsk) gene of *Drosophila melanogaster* differentially affects an operant and a classical learning task. *J Neurosci* **24**: 9745–9751.
- Quinn WG, Harris WA, Benzer S. 1974. Conditioned behavior in *Drosophila melanogaster*. *Proc Natl Acad Sci* **71**: 708–712.
- Sareen P, Wolf R, Heisenberg M. 2011. Attracting the attention of a fly. *Proc Natl Acad Sci* **108**: 7230–7235.
- Seugnet L, Suzuki Y, Stidd R, Shaw PJ. 2009. Aversive phototactic suppression: evaluation of a short-term memory assay in *Drosophila melanogaster*. *Genes Brain Behav* **8**: 377–389.
- Solanki N, Wolf R, Heisenberg M. 2015. Central complex and mushroom bodies mediate novelty choice behavior in *Drosophila*. *J Neurogenetics* **29**: 30–37.
- von Holst E, Mittelstaedt H. 1950. Das Reafferenzprinzip. Wechselwirkungen zwischen Zentralnervensystem und Peripherie. *Naturwissenschaften* **37**: 464–476.
- Waldvogel FM, Fischbach KF. 1991. Plasticity of the landing response of *Drosophila melanogaster*. *J Comp Physiol A* **169**: 323–330.
- Wolf R, Heisenberg M. 1990. Visual control of straight flight in *Drosophila melanogaster*. *J Comp Physiol A* **167**: 269–283.
- Wolf R, Heisenberg M. 1991. Basic organization of operant behavior as revealed in *Drosophila* flight orientation. *J Comp Physiol A* **169**: 699–705.
- Wolf R, Heisenberg M. 1997. Visual space from visual motion: turn integration in tethered flying *Drosophila*. *Learn Mem* **4**: 318–327.
- Yang Z, Bertolucci F, Wolf R, Heisenberg M. 2013. Flies cope with uncontrollable stress by learned helplessness. *Curr Biol* **23**: 799–803.

Received October 29, 2014; accepted in revised form April 9, 2015.